



Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



Policy evaluation during a pandemic[☆]

Brantly Callaway^a, Tong Li^{b,*}

^a Department of Economics, University of Georgia, USA

^b Department of Economics, Vanderbilt University, USA



ARTICLE INFO

Article history:

Received 19 May 2021

Received in revised form 21 October 2022

Accepted 30 March 2023

Available online xxx

JEL classification:

C21

C23

I1

Keywords:

Policy evaluation

Difference-in-differences

Unconfoundedness

Covid-19

Pandemic

Mediators

ABSTRACT

National and local governments have implemented a large number of policies in response to the Covid-19 pandemic. Evaluating the effects of these policies, both on the number of Covid-19 cases as well as on other economic outcomes is a key ingredient for policymakers to be able to determine which policies are most effective as well as the relative costs and benefits of particular policies. In this paper, we consider the relative merits of common identification strategies that exploit variation in the timing of policies across different locations by checking whether the identification strategies are compatible with leading epidemic models in the epidemiology literature. We argue that unconfoundedness type approaches, that condition on the pre-treatment “state” of the pandemic, are likely to be more useful for evaluating policies than difference-in-differences type approaches due to the highly nonlinear spread of cases during a pandemic. For difference-in-differences, we further show that a version of this problem continues to exist even when one is interested in understanding the effect of a policy on other economic outcomes when those outcomes also depend on the number of Covid-19 cases. We propose alternative approaches that are able to circumvent these issues. We apply our proposed approach to study the effect of state level shelter-in-place orders early in the pandemic.

© 2023 Elsevier B.V. All rights reserved.

1. Introduction

There have been a large number of policies implemented in order to decrease the spread of Covid-19. During the early part of the pandemic, the most important of these policies were non-pharmaceutical interventions such as requirements to wear masks, making Covid-19 tests widely available, contact tracing, school closures, lockdowns, and others. These sorts of policies are likely to come with a number of tradeoffs in terms of effectiveness in reducing the spread of Covid-19 as well as their effects on individuals' economic and psychological well-being. Thus, understanding the effects of different policies along a number of dimensions (both effects on number of cases as well as effects on other outcomes) is a key ingredient for researchers, policymakers, and governments to consider when evaluating Covid-19 related policies.

The main way that these policies have been studied by researchers is to compare outcomes in locations that implemented some policy to outcomes in another location that did not implement the policy. Researchers typically exploit

[☆] We thank the editor, associate editor, and two anonymous referees for their insightful comments. We also thank Andrew Goodman-Bacon, Laura Hatfield, Jonathan Roth, Ian Schmutte, Meghan Skira, and seminar participants at Brown University and the 2021 Annual Health Econometrics Workshop for a number of helpful comments. Code for the approaches proposed in the paper are available in the R ppe package which is available at <https://github.com/bcallaway11/pppe>.

* Corresponding author.

E-mail addresses: brantly.callaway@uga.edu (B. Callaway), tong.li@vanderbilt.edu (T. Li).

having access to panel data – data on cases, testing, and economic variables is generally widely available for particular locations over multiple time periods – to try to understand these effects. This sort of setup is very familiar to many researchers in economics, and, with this sort of data availability, an almost default strategy of empirical researchers is to use difference-in-differences. And, indeed, difference-in-differences has been widely used to study the effects of policies in response to Covid-19.

In the current paper, we argue that difference-in-differences has properties that make it relatively less attractive for conducting policy evaluation during a pandemic than it typically would be for most applications in economics. The intuition for our results is that difference-in-differences methods are typically motivated by a two-way fixed effects model for untreated potential outcomes (see, for example, [Blundell and Costa Dias \(2009\)](#)). The key feature of these models is that they include additively separable unit-level unobserved heterogeneity (i.e., a fixed effect). This sort of heterogeneity is very common in applications in economics; a textbook example would be an application on the effect of some policy on individuals' earnings where the researcher is worried that "ability" is unobserved, affects earnings, and is distributed differently between the group of individuals that are affected by the policy and the group of individuals not affected by the policy. In this setup, the unobserved heterogeneity can be differenced out and paths of outcomes for the group of treated units and the group of untreated units can be compared to each other to deliver the effect of the policy.

However, this sort of motivation does not apply in the case of Covid-19. In particular, the main epidemiological models for Covid-19 transmission are highly nonlinear and depend on (i) the number of currently infected individuals in a particular location, (ii) the number of susceptible individuals in a location, and (iii) the transmission properties of Covid-19. In other words, the key challenge for identifying effects of Covid-19 related policies on the number of Covid-19 cases is not that different locations are different in terms of unobserved heterogeneity, but rather that pre-policy differences in the "state of the pandemic" between treated and untreated locations (e.g., differences in the number of Covid-19 cases before the policy is implemented) can lead to substantial differences between locations in how the pandemic would have evolved absent the policy being implemented. These differences can make it difficult to evaluate the effects of Covid-19 related policies. Moreover, these differences are a major concern for evaluating early pandemic policies where different locations' policy choices often depended on the state of the pandemic in those locations.

Another central issue in the economics literature is to understand the effect of Covid-19 related policies on various economic outcomes.¹ Understanding the effects of Covid-19 related policies on economic outcomes is essential in order to understand the costs and benefits of various policies that are aimed at reducing the number of Covid-19 cases. For economics outcomes, we focus on the simple leading case where untreated potential outcomes are generated by a two way fixed effects model that also depends on the current number of Covid-19 cases in a particular location. This setup allows both for the active number of Covid-19 cases to have an effect on the outcome of interest and for the active number of cases to themselves be affected by the policy. In this case, we consider (i) a "standard" version of difference-in-differences that directly compares paths of economic outcomes among treated and untreated locations, and (ii) difference-in-differences when the current number of Covid-19 cases is included as a regressor. We show that, generally, neither approach can deliver the average effect of the treatment on economic outcomes across treated locations. The first strategy breaks down because of differences in the current number of cases across treated and untreated locations. The second strategy breaks down when the policy affects the number of Covid-19 cases (which is the goal of the policy).

We propose alternative approaches that address both sets of issues mentioned above. In particular, for evaluating the effect of policies on Covid-19 cases, we first show that unconfoundedness-type identification strategies (i.e., strategies that compare locations that have the same pre-treatment values of key Covid-19 related variables) do not suffer from the same drawbacks as difference-in-differences approaches. For this case, we propose an estimation strategy that involves estimating (i) the propensity score (i.e., the probability of experiencing the policy conditional on pre-treatment values of Covid-19 related variables) and (ii) an outcome regression for Covid-19 cases in the absence of the policy that is related to the epidemic model. Our approach is doubly robust in the sense that it delivers consistent estimates of policy effects if either the propensity score or outcome regression is correctly specified. This is important as it implies that we can circumvent having to estimate a full structural epidemic model while still delivering estimates of policy effects that are *compatible* with the epidemic model.

Second, for evaluating the effect of Covid-19 related policies on economic outcomes, we propose a two-step approach where the parameters of a two way fixed effects model that additionally allows for current cases to affect the outcome are identified using untreated locations in the first step. Then, in the second step, we recover the path of active cases that treated locations would have experienced on average if they had not been treated (this follows under similar unconfoundedness type arguments used for cumulative cases above). Using these two pieces of information, we are able to construct the average economic outcome that treated locations would have experienced if they had not participated in the policy – and, therefore, the average effect of the policy is identified for treated locations. Unlike other common approaches, this approach allows both for the policy to affect the current number of cases and for the current number of cases to affect untreated potential outcomes.

¹ Throughout the text, we use the term "economic outcomes" but our results apply to any outcome of interest that is outside the epidemic model that we consider in the paper.

We conclude the paper by studying the effect of state-level shelter-in-place orders (SIPOs) on the number of Covid-19 cases and on travel early in the pandemic. These are challenging policies to evaluate because states that implemented these policies also tended to have a large number of Covid-19 cases earlier than states that did not implement this type of policy (or implemented it later). This correlation mechanically leads to larger increases in Covid-19 cases among early treated states relative to untreated and later-treated states. This additionally implies that parallel trends is violated and can lead to difference-in-differences estimates that these policies *increased* the number of Covid-19 cases which is clearly unreasonable. We additionally show that difference-in-differences estimates are very sensitive to minor changes in how they are specified. And, interestingly, despite very different post-policy estimates across different types of DID specifications, none of them are rejected in pre-treatment periods. This suggests that it is not feasible to choose between alternative difference-in-differences specifications based on their performance in pre-treatment periods. Using difference-in-differences, we also sometimes spuriously estimate meaningfully large effects of placebo SIPOs in states that did not actually implement a SIPO. In contrast, we document notably better performance of the unconfoundedness strategy along several dimensions: this approach does not lead to any estimates that SIPOs increased Covid-19 cases, and it does not lead to large or statistically significant effects of placebo SIPOs among states that did not actually implement a SIPO.

Related work

There are a number of recent papers at the intersection of economics, Covid-19, and policy evaluation, and here we only briefly summarize some of the most related ones.

The most related papers to ours are several methodological papers on evaluating Covid-19 related policies. [Allcott et al. \(2020\)](#) propose an event-study regression estimator that is motivated by SIRD models (the same type of model that we consider below) though their approach ends up being substantially different from ours. [Goodman-Bacon and Marcus \(2020\)](#) and [Gauthier \(2021\)](#) discuss two-way fixed effects regressions in the context of evaluating Covid-19 policies. [Chernozhukov et al. \(2021\)](#) propose an alternative approach to evaluate Covid-19 related policies that is motivated by a SIRD model. Their approach is more structural than ours which comes with tradeoffs. For example, their approach can be useful for evaluating counterfactual policies while ours is geared towards evaluating the effects of policies that were actually implemented. On the other hand, our approach generally requires fewer assumptions to evaluate policies that were actually enacted and is set up to be robust to general forms of treatment effect heterogeneity (which is likely to be important in contexts like many early pandemic policies where policies were often implemented at different times in different locations; see [Goodman-Bacon and Marcus \(2020\)](#) for some related discussion). [Alemán et al. \(2020\)](#) provide a way to transform different pandemic “states” across locations and evaluate pandemic-related policies exploiting different locations being in different “stages” of the pandemic. In contrast, our approach essentially compares locations that have been experiencing similar pandemics in pre-treatment periods but where some locations implement the policy and others do not. These are complementary approaches in that our approach tends to require weaker assumptions, but it may not be suitable to study as many early pandemic policies. For example, under certain conditions, the approach in [Alemán et al. \(2020\)](#) can be used to evaluate nationwide pandemic-related policies with common timing across locations by exploiting sub-national variation in the “stage” of the pandemic. This is a setting where our approach would not be viable.

Difference-in-differences has been widely used in empirical work to study the effects of Covid-19 policies. Some examples of papers that consider Covid-19 related policies using difference-in-differences types of identification strategies include [Bartik et al. \(2020\)](#), [Berry et al. \(2021\)](#), [Chetty et al. \(2023\)](#), [Courtemanche et al. \(2020\)](#), [Dave et al. \(2020a,b\)](#), [Gapen et al. \(2020\)](#), [Glaeser et al. \(2021\)](#), [Goolsbee and Syverson \(2021\)](#), [Gupta et al. \(2020\)](#), [Haynes et al. \(2022\)](#), [Hsiang et al. \(2020\)](#), [Juranek and Zoutman \(2021\)](#), [Kong and Prinz \(2020\)](#), [Sears et al. \(2023\)](#), [Wright et al. \(2020\)](#), and [Ziedan et al. \(2020\)](#). [Haber et al. \(2022\)](#) provides a recent survey of empirical strategies that have been used to evaluate Covid-19 related policies.

Finally, on the econometrics side, our paper is related to a large literature on unconfoundedness and difference-in-differences (see [Imbens and Wooldridge \(2009\)](#) for a survey of this literature). More notably, our results on checking the compatibility of structural epidemic models with reduced form identification strategies is broadly similar to a number of papers in econometrics; for example, just in the context of panel data, [Heckman and Robb \(1985\)](#), [Heckman et al. \(1997\)](#), [Athey and Imbens \(2006\)](#), [Blundell and Costa Dias \(2009\)](#), [Chabé-Ferret \(2015\)](#), [Ghanem et al. \(2022\)](#) and [Marx et al. \(2022\)](#) all provide connections between structural models and conditions under which various reduced form approaches, such as difference-in-differences, can be compatible with these models. Our contributions on allowing for infections to both be affected by the policy and to have a direct effect on economic outcomes appear to be conceptually new but related to work on mediation analysis (see [Huber \(2020\)](#) for a summary of this literature) which is also broadly related to work on simultaneous equation models (e.g., [Griliches \(1977\)](#) and [Imbens and Newey \(2009\)](#)). Viewing the “untreated potential” number of Covid-19 infections as a covariate in the model for untreated potential outcomes is related to the idea of difference-in-differences with covariates that can be affected by the treatment; see, [Bonhomme and Sauder \(2011\)](#), [Lechner \(2011\)](#) and [Caetano et al. \(2022\)](#) for related discussion along these lines.

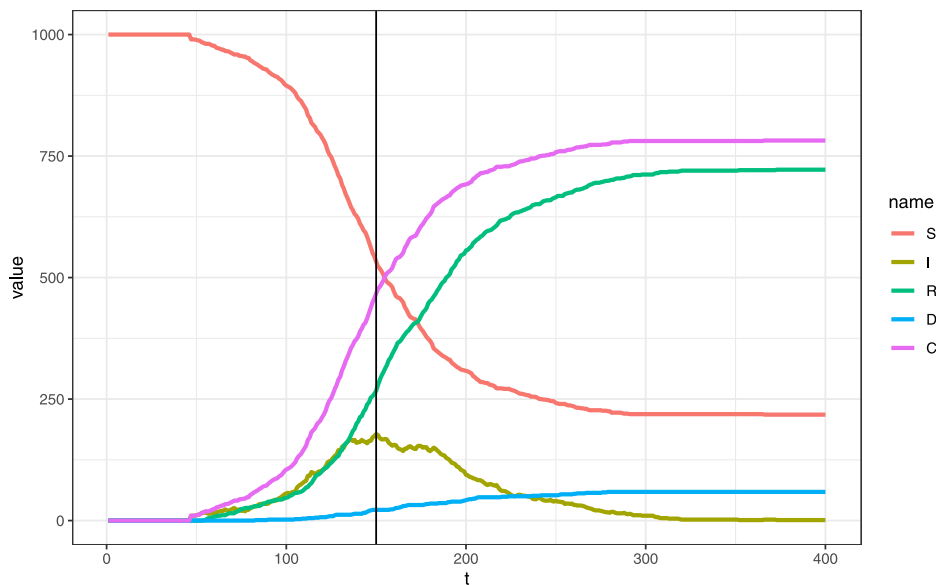


Fig. 1. A simulated SIRD model. *Notes:* The figure shows paths of each variable in a stochastic SIRD model for one (treated) location in the simulated data. The vertical black line is placed at $t = 150$ when the policy was implemented. S stands for the number of susceptible individuals in the population, I stands for the number of currently infected individuals, R stands for the total number of recovered individuals in the population, D stands for the cumulative number of deaths, and C stands for the cumulative number of cases. In this example, the total population is 1000, there are 400 time periods, and the policy has no effect on the pandemic. The specific values for the parameters in this simulation are provided in [Table 5](#) in [Appendix A](#).

2. Motivating examples

To start with, in this section we provide examples of the main types of issues that can confound policy analysis during a pandemic using simulated data from the leading type of epidemic model that has been widely used in the context of Covid-19. [Fig. 1](#) shows the paths of the key variables during a simulated pandemic coming from a stochastic SIRD model. SIRD models categorize individuals in a population into being S-Susceptible, I-Infected, R-Recovered, or D-Dead. We discuss this model in substantially more detail in the next section. The shapes of the path of each variable is typical of a SIRD model. In particular, at some point in time, some small number of cases shows up in a particular location. Then, the number of infections rise in early periods when there are a large number of susceptible individuals in that location combined with an increasing number of currently infected (which also implies contagious). As the number of susceptible decreases (i.e., as infected individuals recover or die), eventually the number of infected individuals decreases. Simultaneously, the cumulative number of cases, number of recovered individuals, and number of deaths all initially grow before eventually leveling off.

In this example, we consider the case where a new policy is implemented in some locations in period 150. For simplicity, we consider the case where the policy has no effect on Covid-19 cases. Locations that participate in the treatment and locations that do not participate in the treatment are alike in all ways except that treated locations tend to experience their first cases earlier than untreated locations. Panel (a) of [Fig. 2](#) shows plots of the average paths of cumulative cases for treated locations and untreated locations in this setup.

Panel (b) of [Fig. 2](#) shows event study-type estimates of the effect of the treatment on the number of cases. To be precise, these are difference-in-differences type estimates where the estimated effect comes from the average change in cases experienced by the treated group of locations relative to the change in cases experienced by the untreated group of locations over the same time periods. Taken at face value, the estimated effects in Panel (b) suggest that the policy decreased the number of Covid-19 cases in treated locations relative to what they would have been in the absence of the policy. However, recall that, in our simulation setup, the policy has no effect on Covid-19 cases. Thus, this example demonstrates that difference-in-differences can perform poorly in the context of trying to evaluate the effect of a policy on the number of Covid-19 cases. The key driver of this poor performance is (i) the nonlinearity of the model for Covid-19 transmission and (ii) differences in the timing of the first cases between locations that participate in the treatment and those that do not. The first of these is an inherent feature of trying to evaluate the effects of policies on Covid-19 cases.

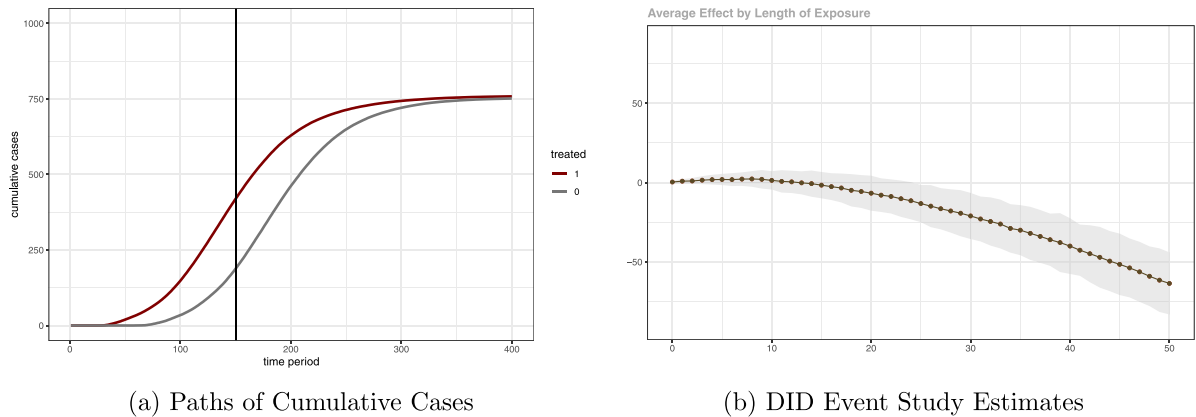


Fig. 2. Simulated policy effects on cumulative cases. *Notes:* Panel (a) plots simulated average paths of cumulative cases among treated and untreated locations. The vertical black line is at $t = 150$ when the policy is implemented for the treated locations. In this example, the policy is constructed so that it has no effect of Covid-19 cases. Panel (b) plots event study type estimates based on a parallel trends assumption for the effect of the policy on the number of cases by length of exposure to the treatment.

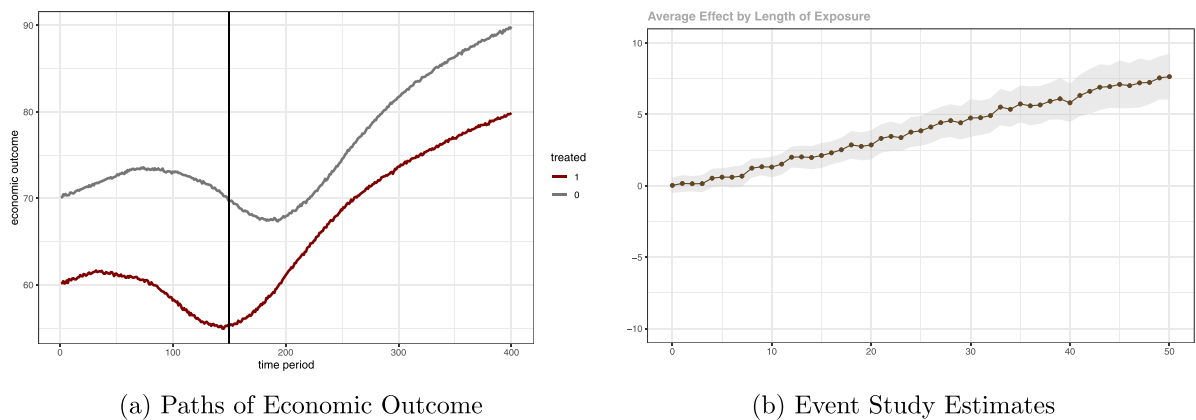


Fig. 3. Simulated policy effects on economic outcome. *Notes:* Panel (a) plots simulated average paths of outcomes among treated and untreated locations. The vertical black line is at $t = 150$ when the policy is implemented for treated locations. In this example, the policy is constructed so that it has no effect on economic outcomes (either directly or through its effect on Covid-19 cases). Panel (b) plots event study type estimates based on a parallel trends assumption for the effect of the policy on the number of cases by length of exposure to the treatment.

For the latter, generally, the bias of difference-in-differences approaches for policy evaluation becomes more severe as the timing of first cases becomes more different between treated and untreated locations.²

Next, we consider the effect of the policy on some economic outcome of interest. Fig. 3 continues with the same simulated policy as above. As above, we consider the case where the policy has no effect on cases or on economic outcomes. However, in this simulation we allow for the economic outcome to depend on the number of active Covid-19 cases in a particular location (here, more active cases tend to decrease the economic outcome), but, otherwise, the economic outcome would follow parallel trends. Panel (a) shows average paths of outcomes for treated and untreated locations in this setup. Panel (b) shows event study type estimates under the assumption of parallel trends. As before, and even in this very simple example, differences in the timing of first cases lead to violations of parallel trends that lead to poor estimates of the effect of the policy on the economic outcome of interest.

Finally, we contrast the poor performance of difference-in-differences in both of these contexts with using an unconfoundedness type strategy to deal with the pandemic-related variables. In particular, when Covid-19 cases is the outcome, we effectively compare locations that were in a similar pandemic “state” in the period right before the policy was implemented. For the economic outcome, we continue to use a version of difference-in-differences, but one that, in the absence of the policy, accounts for economic outcomes depending on the number of Covid-19 infections that would

² Interestingly, the best case for difference-in-differences is when the timing of first cases is the same across treated and untreated locations. However, this is also a case where there is no need to take a time difference at all and one could just make level comparisons of Covid-19 cases across locations.

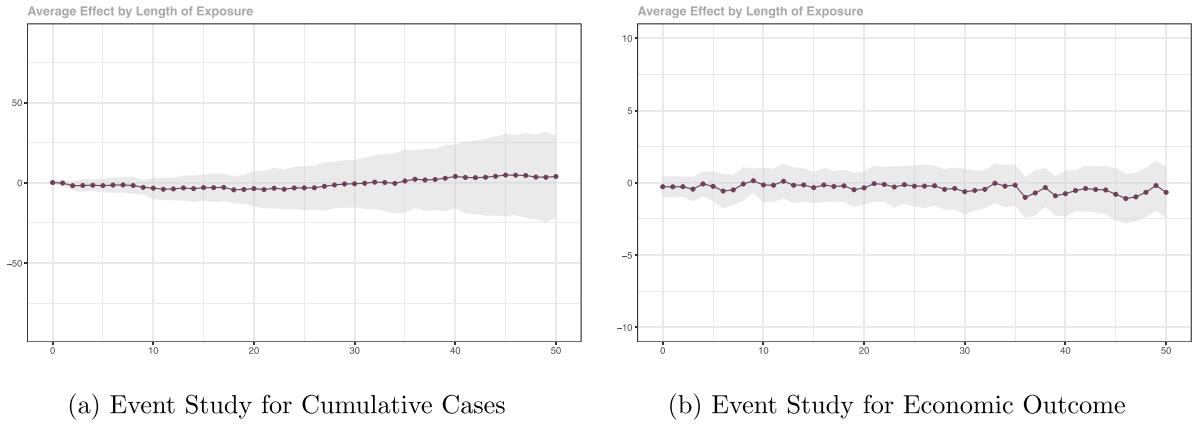


Fig. 4. Estimated policy effects under unconfoundedness. *Notes:* This figure plots event study-type estimates of the effect of a simulated policy on cumulative cases (in Panel (a)) and on an economic outcome (in Panel (b)) using the unconfoundedness-type identification arguments considered in the paper and using the doubly robust estimator discussed below. In this simulation, the policy has no effect on either outcome.

have occurred if the policy had not been implemented using an unconfoundedness strategy (see Section 5 below for more details). Estimates using these approaches are provided Fig. 4. In both cases, these strategies perform notably better at evaluating the effects of the policy.

3. A baseline stochastic SIRD model

In this section, we briefly discuss a stochastic SIRD model which is the workhorse model of epidemic spread in epidemiology and has been used extensively to forecast the spread of Covid-19 cases. SIRD models have a long history in epidemiology – a deterministic version of this kind of model was proposed by Kermack and McKendrick (1927). Stochastic SIRD models are discussed in Allen (2008, 2017) and have been considered by economists in Oka et al. (2021), Fernández-Villaverde and Jones (2022), Ellison (2020) and Acemoglu et al. (2021), among others.

Notation: Let N_l denote the number of individuals in location l . Let \mathcal{T} denote the total number of time periods. The number of susceptible individuals in location l in a particular time period t is denoted by S_{lt} , the number of currently infected individuals in location l at time period t is denoted by I_{lt} , the cumulative number of recovered individuals is denoted by R_{lt} , and the number of cumulative deaths is denoted by δ_{lt} . All individuals in the population are in exactly one of these states at a particular point in time so that

$$N_l = S_{lt} + I_{lt} + R_{lt} + \delta_{lt}$$

in all time periods. Later, we will be interested in the effect of the policy on the cumulative number of cases by time period t , and we denote this variable by C_{lt} and note that $C_{lt} = N_l - S_{lt}$.

In a SIRD model, the paths of all of these variables are governed by some transition equations. The transition equations have the Markov property; i.e., the path of each outcome over time only depends on the “state” of location l in the immediately preceding period. And, in particular, these transition equations are given by

$$\mathbb{E}[I_{lt} | \mathcal{F}_{lt-1}] = (1 - \lambda - \gamma)I_{lt-1} + \beta \frac{I_{lt-1}}{N_l} S_{lt-1} \quad (1)$$

$$\mathbb{E}[R_{lt} | \mathcal{F}_{lt-1}] = R_{lt-1} + \lambda I_{lt-1} \quad (2)$$

$$\mathbb{E}[\delta_{lt} | \mathcal{F}_{lt-1}] = \delta_{lt-1} + \gamma I_{lt-1} \quad (3)$$

$$\mathbb{E}[S_{lt} | \mathcal{F}_{lt-1}] = S_{lt-1} - \beta \frac{I_{lt-1}}{N_l} S_{lt-1} \quad (4)$$

$$\mathbb{E}[C_{lt} | \mathcal{F}_{lt-1}] = C_{lt-1} + \beta \frac{I_{lt-1}}{N_l} S_{lt-1} \quad (5)$$

where for some time period t , we define $\mathcal{F}_{lt} = (S_{lt}, I_{lt}, R_{lt}, \delta_{lt})$, and often refer to this as the “state” of the pandemic in location l in time period t . It is worth considering each of these equations in some more detail. To start with, consider the term $\beta \frac{I_{lt-1}}{N_l} S_{lt-1}$ which shows up in Eqs. (1), (4) and (5). This is the expected number of new cases in time period t conditional on the state of the pandemic in location l in time period $t - 1$. The expected number of new cases from one period to the next depends on three things. First, it depends on I_{lt-1}/N_l which is the fraction of individuals that are infected in period $t - 1$. Holding other things constant, when more individuals are infected, it implies an expected

larger increase in the number of cases. Second, the expected number of new cases depends on the number of susceptible individuals in the population. Intuitively, when there are more susceptible individuals, the number of cases grows more rapidly (other things constant). The spread of a pandemic stops when the number of susceptible becomes small enough which can happen either through “herd immunity” or by decreasing the number of susceptible (for example, through the introduction of a vaccine). Finally, the change in the number of cases depends on the parameter β which is called the infection rate. Most non-pharmaceutical interventions are aimed at changing the infection rate — here, there are two potential benefits: (i) decreasing the infection rate through non-pharmaceutical interventions decreases the total number of cases that need to occur before reaching herd immunity,³ and (ii) if there is a vaccine on the horizon, it also would decrease the total number of cases that occur before herd immunity is reached through the vaccine.

Next, consider Eq. (2). This transition equation says that, on average, the number of total recoveries in location l in time period t (conditional on the state of the pandemic in period $t - 1$) is equal to number of individuals in location l that have already recovered by time period $t - 1$ plus some fraction of infected individuals in period $t - 1$. This fraction is determined by the parameter λ which is the recovery rate from Covid-19. Eq. (3) is the transition equation for deaths. The key parameter is γ which parameterizes the death rate from being infected with Covid-19.

Next, consider Eq. (1). This is the transition equation for active Covid-19 cases. The expected number of infections in period t thus depends on (i) the remaining cases after accounting for recoveries and deaths (this is the first term in Eq. (1)), and (ii) the expected number of new cases (this is the second term in Eq. (1)). Finally, in Eq. (4), the expected number of susceptible individuals is equal to the number of susceptible individuals in time period $t - 1$ minus the expected number of new cases; likewise, the expected number of cumulative cases by time period t is equal to the number of cumulative cases in time period $t - 1$ plus the expected number of new cases.

4 Identification strategies for policy effects on Covid- 19 cases

The previous section presented a basic stochastic SIRD model. This section connects that sort of model with the treatment effects literature and considers the relative merits of difference-in-differences and unconfoundedness strategies for evaluating the effect of a policy on the number of Covid-19 cases.

The strategy of this section is to impose the stochastic SIRD model for untreated potential outcomes and to check if difference-in-differences and/or unconfoundedness are compatible with the stochastic SIRD model. This setup does not place restrictions on how treated potential outcomes (i.e., Covid-19 cases under the policy) are generated. In particular, this is consistent with Covid-19 cases under the policy continuing to follow a stochastic SIRD model but where the values of the parameters potentially change in response to the policy; but it is also more general than that in the sense that there are no substantive restrictions on treated potential outcomes. Perhaps more importantly, this setup also allows for heterogeneous effects of policies across different locations.

Additional treatment effects notation To make the connection with the treatment effects literature, we start by introducing some additional notation. First, we define D_l as a binary variable indicating whether or not location l participated in the treatment. We also define treated and untreated versions of all of the variables in the stochastic SIRD model. In particular, for generic time period t , $S_{lt}(0)$, $I_{lt}(0)$, $R_{lt}(0)$, and $\delta_{lt}(0)$ are the number of susceptible, infected, recovered, and dead individuals in location l in time period t if the policy had not been enacted. We also define $C_{lt}(0)$ as the cumulative number of cases in location l by time period t if the policy had not been enacted. Similarly, we define $S_{lt}(1)$, $I_{lt}(1)$, $R_{lt}(1)$, $\delta_{lt}(1)$, and $C_{lt}(1)$ to be the corresponding treated potential variables; i.e., the values of each of these if the policy had been enacted. Following a large literature on policy evaluation which exploits having access to panel data, we consider the case where the researcher has access to some pre-treatment periods. We suppose that the policy is implemented for treated locations in time period t^* where $1 < t^* \leq \mathcal{T}$.⁴ For random variables indexed by time periods, we define $\Delta X_t := (X_t - X_{t-1})$. Because we are also interested in how policy effects vary across time, some of our arguments involve “long differences” where, for $t_2 > t_1$, we define $\Delta^{(t_1, t_2)} X_t := X_{t_2} - X_{t_1}$. In Appendix A.1, we write the SIRD model given in the previous section in terms of untreated potential outcomes, and we refer to this model as the Stochastic SIRD Model for Untreated Potential Outcomes throughout the remainder of the paper.

Our main interest for this part of the paper is the effect of the policy on the cumulative number of Covid-19 cases. Typically, the main parameter of interest in DID applications (and the parameter that we focus on in the current paper) is the Average Treatment Effect on the Treated (ATT). It is given by

$$ATT_t^C = \mathbb{E}[C_t(1) - C_t(0)|D = 1] \quad (6)$$

where we index the ATT by C to indicate that we are considering the effect of the policy on the cumulative number of cases in time period t . ATT_t^C is the difference between cumulative cases under the policy relative to cumulative cases in the

³ This is also one explanation for repeated “waves” of Covid-19 cases. That is, the infection rate may be temporarily reduced by policy intervention or individual choices but then increases again once these interventions are relaxed.

⁴ In practice, the timing of implementing a particular policy may vary across different locations. Extending our arguments to this case is relatively straightforward, and, therefore, this section considers the case where the policy is implemented at the same time across all treated locations. See Remark 2 below for additional discussion on this point.

absence of the policy on average among locations that participated in the treatment. That this parameter is disaggregated by time period makes it straightforward to report across time periods (as in an event study), but it is also straightforward to, for example, average it across post-treatment time periods in order to report an overall average effect of participating in the treatment.

4.1 Using difference-in-differences to evaluate policy effects on Covid-19 cases

The main underlying motivation for considering a DID approach is when a researcher thinks that untreated potential outcomes are generated from a two-way fixed effects model (see, for example, [Blundell and Costa Dias \(2009\)](#)). These sorts of models are attractive in many applications in economics where there are thought to be important unobserved differences between individuals (or firms, etc.) that are not observed by the researcher. In labor economics, these are often thought of as being unobserved skill; in industrial organization, these may be unobserved differences in productivity across firms; and, in health economics, these may be thought of as proneness to particular health conditions. However, there is an important difference of Covid-19 relative to all of these cases. In general, particular locations do not have time invariant unobservables that make them more or less likely to have a large number of cases; instead, the key differences between locations are (i) the timing of their initial case(s), and (ii) the pandemic response (both in terms of policies and in terms of actions taken by the populations in different locations).⁵

The main result in this section is that there are likely to be major drawbacks to using DID to evaluate the effects of Covid-19 related policies on the number of Covid-19 cases. The two primary reasons for this are (i) the highly nonlinear spread of Covid-19 cases during a pandemic and (ii) that the key difference between locations is the current number of Covid-19 cases rather than some fixed unobserved difference between locations in terms of “proneness” to having a large number of cases.

In this section, we consider whether difference-in-differences approaches are compatible with the stochastic SIRD model presented above. We begin by providing some background on using difference-in-differences to identify the effect of some policy. The key identifying assumption in a DID application is the following parallel trends assumption.

Parallel Trends Assumption. For all $t = 2, \dots, T$

$$\mathbb{E}[\Delta C_t(0)|D = 1] = \mathbb{E}[\Delta C_t(0)|D = 0]$$

The parallel trends assumption says that the path of Covid-19 cases that locations in the treated group would have experienced if they had not participated in the treatment is the same as the path of Covid-19 cases that locations in the untreated group did experience. Invoking this assumption leads to the following estimand for the ATT_t^C for $t \geq t^*$

$$DID_t^C = \mathbb{E}[\Delta^{(t^*-1,t)} C_t | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} C_t | D = 0] \quad (7)$$

where we use the notation DID_t^C to highlight that it may not be equal to ATT_t^C . DID_t^C is equal to the path of Covid-19 cases that treated locations experienced adjusted by the path of Covid-19 cases that untreated location experienced; if the parallel trends assumption holds, then the latter is the path of Covid-19 cases that treated locations would have experienced on average if they had not experienced the policy, and DID_t^C would be equal to ATT_t^C . And, regardless of whether or not the parallel trends assumption holds, DID_t^C is the population quantity for what is estimated in DID applications on Covid-19. Before providing our main result on using difference-in-differences to identify/estimate the effect of a policy on Covid-19 cases, it is also worth mentioning that the primary motivating model for difference-in-differences identification strategies is one where

$$C_{it}(0) = \theta_t + \eta_i + v_{it} \quad (8)$$

where θ_t is a time fixed effect, η_i is location-specific unobserved heterogeneity that can be distributed differently between the treated group and untreated group and v_{it} is an idiosyncratic time varying unobservable. Comparing Eq. (8) to the equation for cumulative Covid-19 cases in Eq. (5), it is immediately clear that these are notably different. In the stochastic SIRD model, the important difference between treated and untreated locations is not unobserved heterogeneity, but rather differences in the current number of Covid-19 cases and the number of susceptible individuals across locations. This immediately provides a suggestive piece of evidence that the parallel trends assumption is unlikely to hold when $C_{it}(0)$ is generated from a stochastic SIRD model.

The next result makes explicit that the parallel trends assumption is generally violated in stochastic SIRD models and provides an expression for the bias resulting from incorrectly imposing the parallel trends assumption in cases where untreated potential outcomes are generated by a stochastic SIRD model.

⁵ One caveat to this is that different locations may have characteristics that are related to the parameters of the SIRD model discussed above. See [Remark 3](#) below for more discussion along these lines.

Theorem 1. *In the stochastic SIRD model discussed above*

$$\mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = d] = \sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = d]|D = d\right]$$

which implies that

- (i) *The parallel trends assumption does not generally hold*
- (ii) *Further, the bias from incorrectly imposing the parallel trends assumption is given by*

$$\begin{aligned} & \left(\mathbb{E}[\Delta^{(t^*-1,t)}C_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}C_t|D = 0] \right) - ATT_t^C \\ &= \left(\sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = 1]|D = 1\right] - \sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = 0]|D = 0\right] \right) \end{aligned}$$

The proof of [Theorem 1](#) is provided in [Appendix B](#). [Theorem 1](#) shows that difference-in-differences generally delivers (potentially severely) biased estimates of the effect of a policy on cumulative Covid-19 cases. It is worth making a few additional comments before proceeding. First, the key reason why the difference-in-differences strategy breaks down is that, in general, the distribution of pandemic related variables immediately before the policy (contained in \mathcal{F}_{t^*-1}) is not the same across treated and untreated locations. Due to the nonlinearity of the SIRD model, this leads to violations of the parallel trends assumption. Second, the sign of the bias cannot generally be determined from these expressions. For example, in [Fig. 2](#) above, difference-in-differences resulted in downward biased estimates of the effect of the policy, but the direction of the bias is sensitive to both (i) timing of first cases in treated and untreated locations, and (ii) the timing of the policy itself (this can be clearly seen in Panel (a) of [Fig. 2](#) where setting the policy at an alternative time period could result in parallel trends being violated in the opposite direction).

Some of the expressions in [Theorem 1](#) seem complicated. One special case of this result that is worth pointing out is when $t = t^*$ (so that we are considering the effect of the policy on Covid-19 cases “on impact”). In that case, the bias from using DID is given by

$$\mathbb{E}\left[\beta \frac{I_{t^*-1}(0)}{N} S_{t^*-1}(0)|D = 1\right] - \mathbb{E}\left[\beta \frac{I_{t^*-1}(0)}{N} S_{t^*-1}(0)|D = 0\right]$$

This bias is the difference between the expected number of new cases that treated locations would have experienced in the absence of the policy relative to the expected number of new cases for untreated locations. And, here, it is straightforward to see key reasons why difference-in-differences can perform poorly: if the joint distribution of currently infected and number of susceptible individuals is different among treated and untreated locations, then they would have experienced a different number of new Covid-19 cases *even if the policy had not been implemented*. In the context of Covid-19, there are some cases where these biases could be substantial. Perhaps the leading example is when the timing of initial Covid-19 cases varied across locations and Covid-19 related policies were implemented earlier in locations that tended to have cases earlier.

4.2 Unconfoundedness in SIRD models

A main alternative to difference-in-differences for evaluating the effects of policies is to assume some version of unconfoundedness. Unconfoundedness means that, after conditioning on some covariates, treatment assignment is as good as randomly assigned.⁶ In other words, in order to identify the effect of some policy on Covid-19 cases, one can compare Covid-19 cases in locations that experienced the treatment to Covid-19 cases in locations that did not participate in the treatment *and had the same characteristics related to the pandemic* as treated locations. In this section, we consider a particular version of unconfoundedness that does not suffer from the same limitations as difference-in-differences for evaluating the effects of policies on the number of Covid-19 cases.

Intuitively, the reason why an unconfoundedness strategy works better for studying policy effects of Covid-19 is that the key differences between locations are the current amount of cases and the current number of susceptible individuals rather than differences in location-specific unobserved heterogeneity. Therefore, conditioning on current cases and the current number of susceptible individuals is sufficient for comparisons of treated and untreated locations to deliver causal effects of policies on Covid-19 cases; while the differencing strategy of difference-in-differences is not able to do the same.

The next result is a main result on the validity of identifying policy effects under the assumption of unconfoundedness. Before stating this result, define the propensity score as

$$p(\mathcal{F}_{t^*-1}) := P(D = 1|\mathcal{F}_{t^*-1})$$

⁶ To be clear on terminology, we use the term “unconfoundedness” rather than, say, “differences-in-differences conditional on pre-treatment characteristics” because, even after taking first differences in the SIRD model, the change in outcomes over time still depends on the lagged outcome. This case ends up reducing to unconfoundedness conditional on the lagged outcome (and other variables); see Footnote 1 in [Abadie \(2005\)](#) for a related discussion.

which is the probability of being treated conditional on pre-treatment characteristics \mathcal{F}_{t^*-1} and make the following assumption

Assumption 1 (Overlap). There exists some $\epsilon > 0$ such that $P(D = 1) > \epsilon$ and $p(\mathcal{F}_{t^*-1}) < 1 - \epsilon$ almost surely.

Assumption 1 is a standard assumption in the treatment effects literature. In the context of Covid-19 related policies, the first part says that there are some locations that participate in the treatment, and the second part says that, for all values of \mathcal{F}_{t^*-1} , one can find untreated locations that have those characteristics. This implies that, for all treated locations, there exists matching untreated locations with the same pre-treatment characteristics. In practice, if this condition is violated, one can identify treatment effects that are local to the region of common support (see, for example, [Crump et al. \(2009\)](#)).

Proposition 1. In the Stochastic SIRD Model for Untreated Potential Outcomes and under [Assumption 1](#), and for any $t \geq t^*$,

$$\mathbb{E}[C_t(0)|\mathcal{F}_{t^*-1}, D = 1] = \mathbb{E}[C_t(0)|\mathcal{F}_{t^*-1}, D = 0]$$

The proof of [Proposition 1](#) is provided in [Appendix B](#). This is an important result and implies that, on average, the unobserved number of cumulative cases that locations that participated in the treatment would have experienced if they had not participated in the treatment is the same as the cumulative number of cases that untreated locations actually did experience among locations that had the same pre-treatment characteristics.

Finally, in this section, we provide an identification result for ATT_t^C which is valid under the SIRD model for Covid-19 cases.

Theorem 2. In the Stochastic SIRD Model for Untreated Potential Outcomes and under [Assumption 1](#), and for any $t \geq t^*$,

$$ATT_t^C = \mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})(C_t - m_{0,t}^C(\mathcal{F}_{t^*-1}))] \quad (9)$$

where

$$\omega(D, \mathcal{F}_{t^*-1}) := \frac{D}{\mathbb{E}[D]} - \frac{\frac{p(\mathcal{F}_{t^*-1})}{1-p(\mathcal{F}_{t^*-1})}(1-D)}{\mathbb{E}\left[\frac{p(\mathcal{F}_{t^*-1})}{1-p(\mathcal{F}_{t^*-1})}(1-D)\right]} \quad \text{and} \quad m_{0,t}^C(\mathcal{F}_{t^*-1}) := \mathbb{E}[C_t|\mathcal{F}_{t^*-1}, D = 0] \quad (10)$$

[Theorem 2](#) says that, under a stochastic SIRD model, we can evaluate the effect of a policy using an unconfoundedness strategy that compares the number of cases in locations that participated in the treatment to the number of cases in locations that did not participate in the treatment and which had the same Covid-19 related characteristics in the period before the policy was implemented.

It is worth making several additional comments related to the result in [Theorem 2](#). First, estimating ATT_t^C from the expression in Eq. (9) involves estimating the propensity score, $p(\mathcal{F}_{t^*-1})$ and the outcome regression $m_{0,t}^C(\mathcal{F}_{t^*-1})$. It is also possible to derive alternative expressions for ATT_t^C that only require either estimating the propensity score (these would be similar to propensity score re-weighting estimators as in [Hirano et al. \(2003\)](#)) or estimating the outcome regression (these would be similar to regression adjustment estimators). However, the expression for ATT_t^C in Eq. (9) possesses the double robustness property.⁷ A main advantage of a doubly robust estimator is that it provides consistent estimates of ATT_t^C if either the propensity score model or the outcome regression model is correctly specified (see, for example, [Bang and Robins \(2005\)](#) and [Słoczyński and Wooldridge \(2018\)](#)). Double robustness is particularly appealing in this context as it enables us to side-step the problem of estimating the full SIRD model and instead involves estimating a model of the treatment assignment process which is both familiar to economists and may be substantially more feasible to do with a simple parametric model. In unreported simulations, we found that imposing flexible parametric models for both the propensity score and the outcome regression performed notably better than either the pure outcome regression approach or the propensity score re-weighting approach.

Second, it is worth briefly mentioning that the weights in Eq. (7) are normalized to have mean one in finite samples. This type of normalized weights is said to be of the Hájek-type ([Hájek \(1971\)](#)) and typically results in estimators with improved finite sample properties relative to its unnormalized counterpart ([Busso et al. \(2014\)](#)). Finally, we provide the asymptotic properties of our estimator in the Supplementary Appendix. In order to conduct inference, we use a multiplier bootstrap procedure that involves perturbing the influence function of the estimator of ATT_t^C ; we also discuss how to conduct uniform inference across different time periods to account for multiple testing. These results primarily follow from recent results on doubly robust estimators with Hájek-type weights in [Sant'Anna and Zhao \(2020\)](#).

Remark 1. The results in this section have focused on the effect of a policy on the number of cumulative cases. However, the same sorts of arguments apply to other possible variables of interest such as the current number of infections. For example, the proof of [Proposition 1](#) additionally covers the other variables in the epidemic model, and we use similar arguments as in [Theorem 2](#) but for the current number of infections in the next section.

⁷ For completeness, we provide a proof in the Supplementary Appendix, but the arguments follow along the same lines as arguments for existing doubly robust estimators under unconfoundedness.

Remark 2. The identification arguments in this section have been for the case where the timing of the policy does not vary across different locations. However, it is straightforward to extend these arguments to the case where the timing of the policy varies across locations (as is the case in our application). In this case, one can think of our identification arguments holding specifically for each “group” where a group is defined by the time period when a location first becomes treated. In this case, instead of identifying ATT-type parameters, one would identify group-time average treatment effects along the lines of [Callaway and Sant’Anna \(2021\)](#) and can follow their approach to aggregating these sorts of parameters into an overall average effect of participating in the treatment or into an event study type result. We provide a complete discussion of this extension in the Supplementary Appendix. There is variation in treatment timing in the application that we consider below, and this is the approach that we follow in the application.

Remark 3. In the Supplementary Appendix, we consider a number of extensions to the Stochastic SIRD Model for Untreated Potential Outcomes and pay particular attention to which sorts of extensions are or are not compatible with the unconfoundedness condition discussed in this section. The types of extensions that we consider allow for the parameters of the SIRD model to vary by location and time periods; for example, instead of the infection rate β being constant across locations and time periods, we instead denote the infection rate by β_{lt} and consider what sort of additional structure on β_{lt} is compatible with unconfoundedness. First, we show that unconfoundedness is generally compatible with SIRD model parameters varying over time; e.g., $\beta_{lt} = \beta_t$. This can be empirically relevant if, for example early in the pandemic, the availability of masks changed over time or when there is common time-varying information about how Covid-19 is transmitted. On the other hand, we also show that the unconfoundedness approach is not compatible with model parameters that are specific to particular locations. One example is when $\beta_{lt} = \exp(\theta_t + \theta_l)$ where $\exp(\cdot)$ enforces that the infection rate should be positive and θ_t is a time fixed effect and θ_l is a location fixed effect in the infection rate. An empirically relevant exception to this limitation is when the location-specific fixed effects can be accounted for by observed location-specific characteristics. For example, arguably the most important reasons for location-specific variation in infection rates are likely to be differences in population density, age and/or income distribution, and demographic characteristics – all of which are generally observed (see, [Chernozhukov et al. \(2021\)](#) for some related discussion along these lines).

5 Identification strategies for policy effects on economic outcomes

Another interest of economists is studying the effect of Covid-19 related policies on other (particularly economic) outcomes. This is likely to be useful for thinking about a cost–benefit analysis of particular policies. Relative to textbook versions of difference-in-differences, what is different in this section is that we allow for economic outcomes to depend on the current number of Covid-19 cases in a particular location.⁸ We denote the economic outcome of interest by Y_{lt} which is the observed economic outcome for location l in time period t . We also define treated potential outcomes, $Y_{lt}(1)$, and untreated potential outcomes, $Y_{lt}(0)$, and note that $Y_{lt} = D_l Y_{lt}(1) + (1 - D_l) Y_{lt}(0)$. The target parameter in this section is given by

$$ATT_t^Y = \mathbb{E}[Y_t(1) - Y_t(0)|D = 1]$$

which is the difference between treated potential outcomes and untreated potential outcomes on average, in time period t , and among treated locations. As discussed above, difference-in-differences is closely related to two-way fixed effects models for untreated potential outcomes, and, in this section, we consider the following model for untreated potential outcomes

$$Y_{lt}(0) = \tau_t + \xi_l + \alpha I_{lt}(0) + v_{lt} \quad (11)$$

where τ_t is a common macro shock. For economic outcomes, there is clear evidence of common macroeconomic shocks which can be motivated by, for example, common information about the health risks of Covid-19 across locations. ξ_l is a location-specific fixed effect allowing for time-invariant location-specific differences in economic outcomes, and v_{lt} are idiosyncratic, time varying unobservables.

What is different about this model from standard DID is the term involving $I_{lt}(0)$ where $I_{lt}(0)$ is the number of Covid-19 cases in location l in time period t if the policy were not implemented. It is likely to be very important to include this sort of term during the pandemic as it allows for economic outcomes to depend on the local spread of cases. In particular, this allows for current cases to directly affect outcomes as well as individuals and/or firms taking more Covid-19 related precautions when the number of local cases is high.

In this section, we propose an approach that is able to deliver consistent estimates of ATT_t^Y in the case when policies can have an effect on current Covid-19 cases and current Covid-19 cases can, in turn, have an effect on the outcome of interest. Throughout this section, we contrast our suggested approach with two very common DID-type approaches. First, we consider the case where a researcher compares the path of outcomes of treated locations to the path of outcomes

⁸ As above, because the target parameter is an ATT-type parameter, the setup in this section does not require assumptions on how treated potential outcomes are generated and, therefore, the discussion about the effect of current cases and SIRD models in this section need only apply for untreated potential outcomes.

among untreated locations without accounting for the current number of Covid-19 cases. Throughout this section, we refer to this case as “standard DID”. Second, we consider a version of DID that includes the number of cases as a regressor. Throughout this section, we refer to this case as “regression DID”. We show that both of these approaches generally deliver biased estimates of ATT_t^Y under the model in Eq. (11). In the standard DID case, biased estimates arise because the researcher does not account for current cases in a particular location having a direct effect on outcomes. In the regression DID case, biased estimates arise because the approach does not accommodate the possibility that the policy has an effect on the current number of cases (which in turn has an effect on outcomes).

In light of this discussion, we propose an alternative approach that simultaneously addresses both of these issues. We call our approach “adjusted regression DID”. Our idea is to include an adjustment term that accounts for the possibility that the policy affects the current number of cases. This adjustment term is closely related to the arguments in the previous section; in particular, we can recover an estimate of the number of active cases that a treated location would experience in a particular time period by recovering the number of active cases in untreated locations with similar pre-treatment pandemic-related characteristics.

Before stating the main result in this section, it is helpful to notice that, in the model in Eq. (11),

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D = d] = \tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = d] \quad (12)$$

where we define $\tilde{\tau}_t := (\tau_t - \tau_{t^*-1})$. Also note that $\tilde{\tau}_t$ and α are both identified using the untreated group (in that case, untreated potential outcomes and untreated potential active cases are observed in all time periods which implies that the parameters are identified as this amounts to a simple linear regression of $\Delta^{(t^*-1,t)}Y_t$ on $\Delta^{(t^*-1,t)}I_t$ using untreated locations).

Next, to fix ideas, under standard DID, the estimator of ATT_t^Y is the sample analogue of

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 0]$$

Likewise, under regression DID, the estimator of ATT_t^Y is the sample analogue of

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - (\tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D = 1])$$

Including covariates in this sort of way is a common strategy⁹ and would amount to comparing paths of outcomes for treated and untreated locations that experienced the same change in cases over time.

The next result provides an alternative identification result for ATT_t^Y as well as results for the bias of standard DID and regression DID.

Theorem 3. *In the model for untreated potential outcomes in Eq. (11) and the Stochastic SIRD Model for Untreated Potential Outcomes and under Assumption 1, and for any $t \geq t^*$,*

$$\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1] = \mathbb{E}\left[\frac{D}{\mathbb{E}[D]} \Delta^{(t^*-1,t)}I_t - \omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^I(\mathcal{F}_{t^*-1}))\right] \quad (13)$$

where $m_{0,t}^I(\mathcal{F}_{t^*-1}) := \mathbb{E}[I_t|\mathcal{F}_{t^*-1}, D = 0]$ and ω are the same weights as in Theorem 2. Moreover,

$$ATT_t^Y = \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - \left\{\tilde{\tau}_t + \alpha \left(\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1]\right)\right\} \quad (14)$$

and all the terms on the RHS of the expression for ATT_t^Y are identified. In addition, standard difference-in-differences is biased with bias given by

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 0] - ATT_t^Y = \alpha(\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 0])$$

and the bias of regression DID (i.e., including current cases as a covariate) is given by

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - (\tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D = 1]) - ATT_t^Y = -\alpha ATT_t^I$$

where $ATT_t^I := \mathbb{E}[I_t(1) - I_t(0)|D = 1]$.

It is worth sketching the arguments underlying the result in Theorem 3. To start with, notice that

$$ATT_t^Y = \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D = 1]$$

The bias of standard DID arises from (incorrectly) setting $\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D = 1] = \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 0]$. In general, this sort of substitution is not appropriate because the path of outcomes that treated locations would have experienced in the absence of participating in the treatment depends on the path of active cases (which is not accounted for here).

⁹ Notice that this estimand is similar in spirit, though not exactly the same, as two way fixed effects regressions that include a treatment dummy variable along with other time varying covariates. Besides the issues pointed out in this section (related to the covariates), those sorts of regressions do not generally deliver an interpretable treatment effect parameter in the case with multiple time periods and variation in treatment timing (see, for example, Goodman-Bacon (2021)). The estimand mentioned above avoids the issues related to multiple periods and variation in treatment timing but, as we point in this section, still suffers from issues stemming from actual cases in treated locations not being equal to what cases would have been if the policy had not been implemented.

Next, based on the model in Eq. (11), it follows from Eq. (12) that

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D = 1] = \tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1] \quad (15)$$

The bias of regression DID (that directly includes current cases as a covariate) comes from (incorrectly) setting $\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1] = \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D = 1]$. This strategy is also not generally appropriate because the policy can change (and is likely targeted at changing) the path of active cases.

By contrast, our approach uses the expression in Eq. (13) for $\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1]$. This expression takes the observed path of active cases and subtracts from it the effect of the policy on active cases (which is the term $\mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^l(\mathcal{F}_{t^*-1}))]$) and holds under the Stochastic SIRD Model for Untreated Potential Outcomes). Notice that this term is analogous to the expression for the effect of the policy on cumulative cases in Theorem 2 in the previous section. The difference between the observed path of active cases among treated locations and the effect of the policy on active cases recovers the path of active cases that would have occurred if the policy had not been implemented. Given this expression, it can be plugged into Eq. (15) to recover the path of untreated potential outcomes and, hence, to recover ATT_t^Y .

Estimation: The above discussion suggests the following estimation strategy:

1. Run a regression of $\Delta^{(t^*-1,t)}Y_{it}$ on $\Delta^{(t^*-1,t)}I_{it}$ using untreated locations in order to estimate $\tilde{\tau}_t$ and α .
2. Estimate the propensity score, $p(\mathcal{F}_{t^*-1})$, using the entire sample. Plug these estimates into the sample analogue of Eq. (10) in order to estimate the weights ω . Estimate $m_{0,t}^l(\mathcal{F}_{t^*-1})$ using untreated locations. Plug in the estimates of ω and $m_{0,t}^l$ into the sample analogue of Eq. (13) to compute an estimate of $\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1]$.
3. Plug in the preliminary estimators in Steps 1 and 2 to estimate ATT_t^Y directly using the expression in Theorem 3.

In the Supplementary Appendix, we provide the asymptotic distribution of our estimator of ATT_t^Y . The estimation procedure involves several steps, but each step is parametric and the limiting distribution of the estimate of ATT_t^Y can be obtained following well-known arguments about multiple step estimation procedures that account for estimation effects of each step. In particular, the term $\mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^l(\mathcal{F}_{t^*-1}))]$ can be handled using exactly the same arguments as in the previous section. The other steps in the estimation procedure only involve either running simple parametric regressions or directly calculating averages and are therefore straightforward to account for. As earlier, in practice, we use the multiplier bootstrap to conduct inference and discuss how to conduct uniform inference across different time periods.

Remark 4. In general, it is not possible to sign the bias from using standard DID or regression DID (as an extreme example, over long enough time horizons, a policy that is effective at slowing the spread of Covid-19 could lead to higher current infections if untreated locations reach herd immunity). That said, over relatively short time horizons, it is possible to get a sense of the *likely* directions of bias. In particular, suppose that (i) $\alpha < 0$ so that more current cases lead to lower economic outcomes, (ii) that the time horizon is short and the policy decreases the number of active cases over a short horizon, (iii) that the pandemic is in its early stages and that treated locations tend to have earlier arrival times of their first cases (so that, in the absence of participating in the treatment, treated locations would have experienced larger increases in the number of active cases than untreated locations), and (iv) the policy has a negative effect on economic outcomes. In this case, both standard DID and regression DID (that adjusts for the actual number of cases) will both overstate the magnitude of the effect of the policy.

6 Monte Carlo simulations

In this section, we provide some Monte Carlo simulations to demonstrate the performance of the main estimation strategies considered in the paper. To begin with, we consider estimating the effect of a policy on cumulative Covid-19 cases. In order to generate the data, we consider the case where untreated potential outcomes are generated by the Stochastic SIRD Model for Untreated Potential Outcomes. The values for the main parameters in the SIRD model are provided in Table 5 in Appendix A. We also suppose that the policy has no effect on the pandemic so that all treatment effects are equal to 0.

Throughout this section, we consider the case where there are 250 locations (we vary this number in a few cases), where the probability of a location being treated is equal to 0.5, and where there are 1000 individuals in each location. We report bias, root mean squared error, and rejection probabilities for $H_0 : ATT = 0$ for the average effect of the policy across the first 50 post-treatment time periods (i.e., we compute event-study type estimates for 50 periods following the treatment, average them across event time to get an overall average treatment effect parameter, and compute the properties of this estimator). To implement our doubly robust estimator, we include a third order polynomial (also including all interactions) in the pre-treatment number of infected individuals and pre-treatment number of susceptible individuals both for the outcome regression and for the propensity score. Across simulations, we primarily focus on varying the timing of initial Covid-19 cases among treated and untreated locations, and on varying the treatment timing across treated and untreated locations.

The results for our first set of simulations are provided in Table 1. The high level takeaway from this table is that the unconfoundedness approach uniformly appears to perform better than difference-in-differences. Difference-in-differences is severely biased when the timing of initial cases is different between treated and untreated locations (this

Table 1
Monte Carlo simulations for cumulative cases.

Policy Time	λ_D	λ_U	Unconfoundedness			DID		
			Bias	RMSE	Rej. Prob.	Bias	RMSE	Rej. Prob.
Vary treated first case timing								
150	40	60	0.009	0.478	0.039	−3.044	4.169	0.162
150	60	60	0.008	0.582	0.044	0.031	3.233	0.055
150	80	60	−0.012	0.750	0.065	2.931	4.542	0.153
Vary policy timing								
75	40	80	0.034	0.803	0.036	−12.829	14.416	0.469
150	40	80	0.034	0.428	0.024	−5.593	6.464	0.438
225	40	80	0.047	0.196	0.031	−1.133	1.389	0.323
Vary number of locations, $n = 1000$								
150	40	80	0.031	0.194	0.044	−5.680	5.895	0.951

Notes: The table provides Monte Carlo simulations for cumulative cases using the unconfoundedness approach suggested in the paper as well as difference-in-differences and with the simulation parameters discussed in the text. The column labeled “Policy Time” indicates the timing when the policy is implemented among treated locations; λ_D and λ_U are the mean timing of the first case for treated and untreated locations, respectively. The other columns report the bias, root mean squared error (RMSE), and rejection probabilities for each simulation setup.

Table 2
Monte Carlo simulations for economic outcomes.

Policy Time	λ_D	λ_U	Adjusted regression DID			Standard DID		
			Bias	RMSE	Rej. Prob.	Bias	RMSE	Rej. Prob.
n = 250								
150	40	60	0.000	0.127	0.048	−0.134	0.227	0.092
150	60	60	0.001	0.132	0.048	0.002	0.193	0.043
150	80	60	−0.015	0.132	0.049	0.110	0.230	0.081
n = 1000								
150	40	60	0.003	0.066	0.055	−0.129	0.159	0.263
150	60	60	0.005	0.067	0.045	0.005	0.098	0.051
150	80	60	0.001	0.071	0.068	0.127	0.165	0.240

Notes: The table provides Monte Carlo simulations for economic outcomes using the adjusted regression DID approach suggested in the paper as well as standard DID and with the simulation parameters discussed in the text. The column labeled “Policy Time” indicates the timing when the policy is implemented among treated locations; λ_D and λ_U are the mean timing of the first case for treated and untreated locations, respectively. The other columns report the bias, root mean squared error (RMSE), and rejection probabilities for each simulation setup.

is in line with our earlier discussion). The magnitude of the bias of difference-in-differences is also sensitive to the timing of the policy (this holds because the direction/magnitude of violations of parallel trends depends on the shape of pandemic related variables which are, in turn, dependent on how long ago the pandemic started). Across simulations, difference-in-differences also tends to over-reject.

On the other hand, the doubly robust unconfoundedness approach performs much better with good performance across each specification. Interestingly, even in the case where the first cases show up, on average, at the same time across treated and untreated locations (in this case, as expected, DID appears to be unbiased), the unconfoundedness approach suggested in the paper has notably smaller root mean squared error.

Next, we provide analogous results but for the effect of the policy on economic outcomes. For this part, we generate untreated potential outcomes according to Eq. (11). We set $\alpha = -0.1$, $\tau_t = (50 + 20 \times t/\mathcal{T})$, and we set $\eta|D = d \sim N(\mu_d, 1)$ where $\mu_d = 20 - 10d$ for $d \in \{0, 1\}$. We also set the parameters for the pandemic related variables as in the baseline specification discussed above.

These results are provide in Table 2 where we vary the timing of initial cases as well as the number of locations across simulations. As in the previous case, the approach suggested in the paper (adjusted regression DID) performs well uniformly across DGPs. By contrast, standard DID performs less well particularly in cases where the timing of initial cases is systematically different across treated and untreated locations.

7 Application: Effects of shelter-in-place orders on Covid-19 cases and travel

To conclude the paper, we apply our approach to study the effect of shelter-in-place orders (SIPOs) on Covid-19 cases and travel. We start by using state-level data to evaluate the effect of SIPOs on Covid-19 cases. This approach is broadly similar to a number of other papers including Berry et al. (2021), Bendavid et al. (2021), Courtemanche et al. (2020), Dave et al. (2021), Dave et al. (2020b), Hsiang et al. (2020), and Sears et al. (2023), among others. We consider a number of variations of difference-in-differences estimation strategies in this context as well as implementing the unconfoundedness approach discussed above. Importantly, we document substantial differences in important pre-treatment characteristics such as the number of Covid-19 cases between states that were early adopters of SIPOs, late adopters of SIPOs, or never

implemented a SIPO. As emphasized above, this suggests that the parallel trends assumption underlying the DID approach is likely to be violated. We find that DID approaches can, in some cases, lead to unreasonable estimates that SIPOs *increased* Covid-19 cases. Another of our main findings is that the DID approach is quite sensitive to seemingly minor modifications to the specification such as using the logarithm or level of the outcome or whether or not one includes location-specific linear trends.

For any approach using state-level data, there are some important challenges. One of these is that a number of additional Covid-19 related policies such as emergency declarations, school closures, and business closures were often implemented around the same time. Moreover, there is variation across states both in terms of exactly which policies were implemented as well as the timing of these policies; we discuss some other challenges below as well. Therefore, as a second step, we use county-level data and provide difference-in-differences and unconfoundedness estimates of policy effects in states where a SIPO was implemented relative to bordering states that never implemented a SIPO among states that are similar both in terms of their pre-treatment pandemic-related characteristics (such as having experienced a similar number of Covid-19 cases) and in terms of the mix and timing of other policies that were implemented.

In general, using variations of DID, we tend to find a hard-to-interpret mix of policy effects that include estimates indicating large reductions in Covid-19 cases due to SIPOs, no effect of SIPOs, or even relatively large *increases* in Covid-19 cases due to SIPOs. This contrasts with our results using unconfoundedness which are more consistent across different state-level policies and where we tend to find relatively smaller reductions in Covid-19 cases due to SIPOs. Finally, we also use the county-level data to study the effects of SIPOs on travel. In line with the literature (e.g., Goolsbee and Syverson (2021)), we tend to find relatively small reductions in travel due to SIPOs.

7.1 State-level results

Our first set of results come from analyzing state-level data. We consider a period early in the pandemic – March 10, 2020 to May 1, 2020 – when a large number of states implemented shelter-in-place orders.

Data

We follow Dave et al. (2021) in terms of definitions of shelter-in-place orders and the timing of implementation across states. In order to facilitate estimating conditional treatment assignment probabilities, we assign states into “groups” on the basis of the timing when they adopted a SIPO. And, in particular, we assign states that adopted a SIPO within a five day window, starting on March 19, into the same group. For example, California was the first state to implement a shelter-in-place order on March 19; Illinois and New Jersey followed on March 21; New York on March 22; and Connecticut, Louisiana, Oregon, and Washington on March 23. These form a group of states that we refer to as the March 19 group. Sixteen other states adopted shelter-in-place orders between March 24 and March 28 and form the group that we refer to as the March 24 group. We include four such groups total as well as an untreated group of ten states that did not adopt a shelter-in-place order over the time period that we consider.

Next, we obtained data on state-level Covid-19 cases and testing from the Centers for Disease Control COVID Data Tracker (<https://covid.cdc.gov/covid-data-tracker/>). We also use 2019 state-level populations from the Census Bureau. In order to deal with heterogeneity in terms of state populations, we use versions of pandemic related variables in terms of their number per thousand individuals in a particular state (e.g., cumulative cases per thousand individuals). In terms of the SIRD model, this amounts to dividing all variables by N_i and multiplying by one thousand; this transformation is compatible with the SIRD model. We construct the current number of active Covid-19 cases (and therefore contagious individuals) as the total number of newly reported cases over the past seven days; as for the other pandemic related variables, we use the number of current cases per thousand individuals in a state. Finally, we use travel data from Google’s Covid-19 Community Mobility Report (<https://www.google.com/covid19/mobility>). We focus on state-level retail and recreation travel (these are aggregated together) which is reported as a percentage change relative to pre-Covid travel baselines.

Summary statistics for the data that we use are provided in Table 3. There are some things that are immediately notable from the summary statistics. First, early in the pandemic, the number of cases were substantially different for states that adopted shelter-in-place orders earlier relative to states that adopted them later or that did not adopt them at all. This immediately suggests that it will be challenging for DID to perform well at evaluating the effect of shelter-in-place orders on the number of cases. In addition, notice that early treated states (particularly, the March 19 group) experienced very large increases in their number of cases relative to later- and never-adopters of the policy. Finally, the second panel of the table shows changes in retail and recreation trips. The most notable feature of this part of the table is that there were large decreases in travel across all states regardless of their shelter-in-place policies.

Results

Our first set of results come from using state-level data and several different types of difference-in-differences estimation strategies. In Fig. 5, we use two-way fixed effects (TWFE) event study regressions of the form

$$C_{it} = \theta_t + \eta_i + \sum_{e=-(T-1)}^{-2} \beta_e D_{it}^e + \sum_{e=0}^{T-1} \beta_e D_{it}^e + v_{it} \quad (16)$$

Table 3
Summary statistics for state-level data.

	Group				
	Untreated	Mar 19	Mar 24	Mar 29	Apr 3
Group size	10	8	16	10	6
pop.(millions)	2.9	12.6	3.7	8.8	8.5
<u>Cases:</u>					
Mar 14	12.2	15.0	15.3	3.1	17.9
Mar 21	52.3	114.6	82.4	27.9	89.7
Mar 28	190.0	573.3	257.3	144.0	267.8
Apr 4	440.8	1609.6	566.0	392.3	555.6
Apr 11	807.0	2742.6	944.1	739.7	902.0
Apr 18	1285.1	3696.1	1324.1	1075.2	1223.4
Apr 25	1917.6	4701.9	1765.3	1475.7	1567.1
<u>Travel</u>					
Mar 14	−9.8	−10.5	−6.9	−5.7	−1.3
Mar 21	−39.4	−43.1	−40.4	−38.7	−34.0
Mar 28	−44.5	−53.5	−52.9	−45.2	−39.8
Apr 4	−45.7	−52.6	−49.0	−47.9	−46.5
Apr 11	−42.2	−49.8	−45.8	−42.7	−41.3
Apr 18	−36.9	−52.0	−43.8	−41.9	−36.3
Apr 25	−33.2	−47.5	−38.5	−39.2	−34.2

Notes: The table provides summary statistics for the state-level data used in the application. A state's "group" is defined by the time period when it became treated rounded to the nearest 5th day (see text for a detailed explanation). Cases are reported as the number per million individuals in a state and averaged within the corresponding "group." Travel data is reported as the percentage change in retail and recreation trips relative to pre-Covid baseline.

where D_{lt}^e is a binary variable that is equal to 1 for location l in time period t if that location has been treated for exactly e periods in period t and is otherwise equal to 0. These sorts of regressions have been widely used in work evaluating the effects of Covid-19 related policies. The panels in Fig. 5 differ along two dimensions: first, whether the outcome is the logarithm or the level of the number of Covid-19 cases; and second, whether or not the specification additionally includes a location-specific linear time trend. All of these specifications are common in the literature, and, in particular, the specification where the outcomes is in logarithms and includes a location-specific linear trend is a main specification in Dave et al. (2021).

Fig. 5 highlights that the qualitative conclusion as to whether or not SIPOs affected the number of Covid-19 cases is highly sensitive to functional form assumptions made by the researcher. First, when the outcome is the logarithm of the number of Covid-19 cases and the specification includes a location-specific linear time trend, then the estimates indicate a large reduction in Covid-19 cases due to shelter-in-place policies (see panel (d) of Fig. 5). The results in panel (c), which also include a linear trend but where the outcome is the level of the number of cases, also indicate that SIPOs may have reduced the number of Covid-19 cases (these results are closer to zero and only marginally statistically significant). On the other hand, the results in panels (a) and (b) of Fig. 5, neither of which include a location-specific linear trend, are much different and suggest that SIPOs led to an *increase* in the number of Covid-19 cases. It seems very hard to rationalize these sorts of results; in particular, it would seem that shelter-in-place orders could either have no effect or decrease Covid-19 cases, but it is difficult to see how they could increase cases. However, even from the summary statistics, one can see that DID estimates are likely to be positive as early policy adopters were tending to experience larger increases in cases. In our view, a better explanation for these results is that Covid-19 was more prevalent earlier in locations that were early policy adopters and that the strong, early exponential growth of Covid-19 cases overwhelms any reduction in the infection rate due to the policy. It is exactly in this case where DID would be susceptible to attributing faster growth in Covid-19 cases to the policy rather than to simply a larger number of early cases in treated locations.

Another noteworthy feature of Fig. 5 is that none of the four specifications used in the figure result in either large or statistically significant violations of the underlying parallel trends assumption in pre-treatment periods (i.e., estimates different from 0 for $e < 0$).¹⁰ This suggests that it is not possible to use a purely data-driven/reduced form model selection procedure to infer that one specification is likely to perform better than the others, despite the choice of the model largely driving the results.¹¹ Finally, we should emphasize that *none* of the TWFE event study specifications discussed here are compatible with the SIRD model that we have discussed in the paper.

¹⁰ To be more specific, there are 23 pre-treatment estimates reported in each panel. No pre-treatment estimates are statistically significant at the 5% level in panels (a), (b), or (c). One estimate is statistically significant at the 5% level in panel (d); this occurs for $e = -3$ and the p -value is 0.037.

¹¹ In one sense, it is somewhat surprising that parallel trends cannot be rejected in pre-treatment periods for any of the estimation strategies considered here. The Stochastic SIRD Model for Untreated Potential Outcomes implies that, for all of the models considered here, the parallel trends assumption is also violated in pre-treatment periods. Instead, the results here indicate that we are not able to *detect* violations of parallel trends in pre-treatment periods (even though parallel trends is probably actually violated). That we are not able to detect violations of parallel trends is not altogether surprising as pre-tests are often under-powered (Roth (2022)), and this is likely to be an acute issue early in the pandemic when the number of Covid-19 cases is very small in most pre-treatment periods.

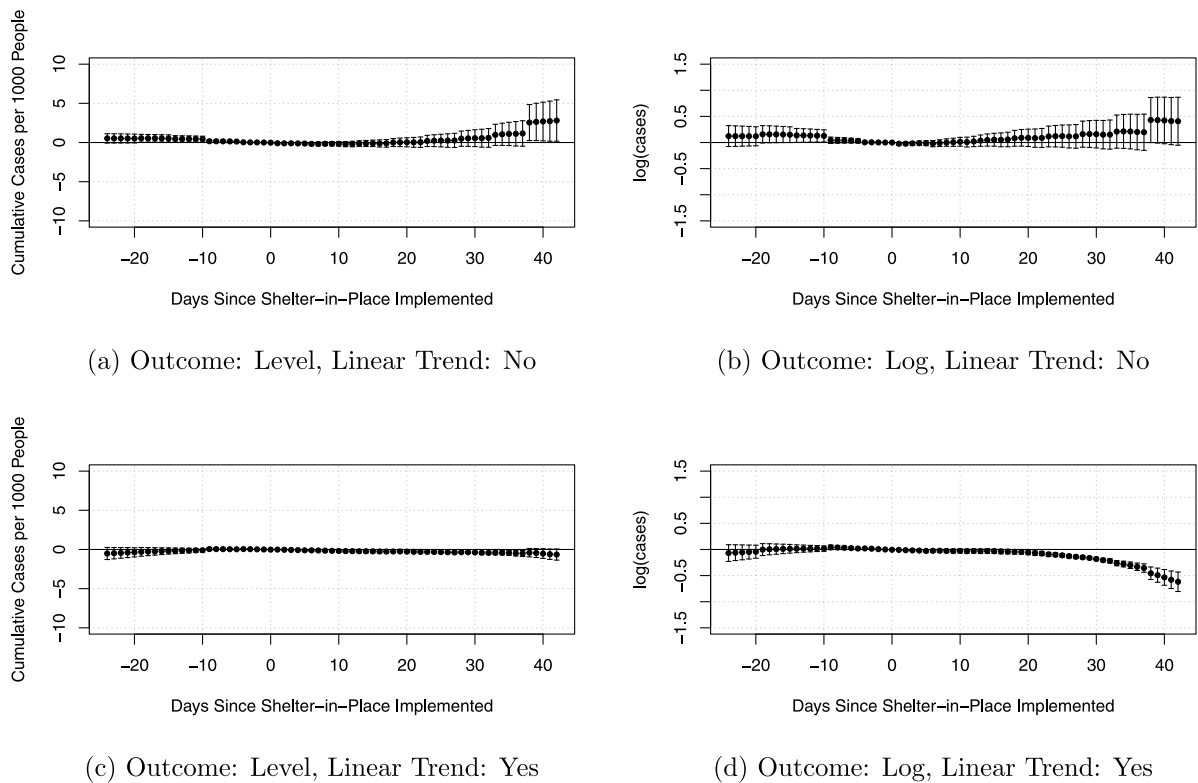


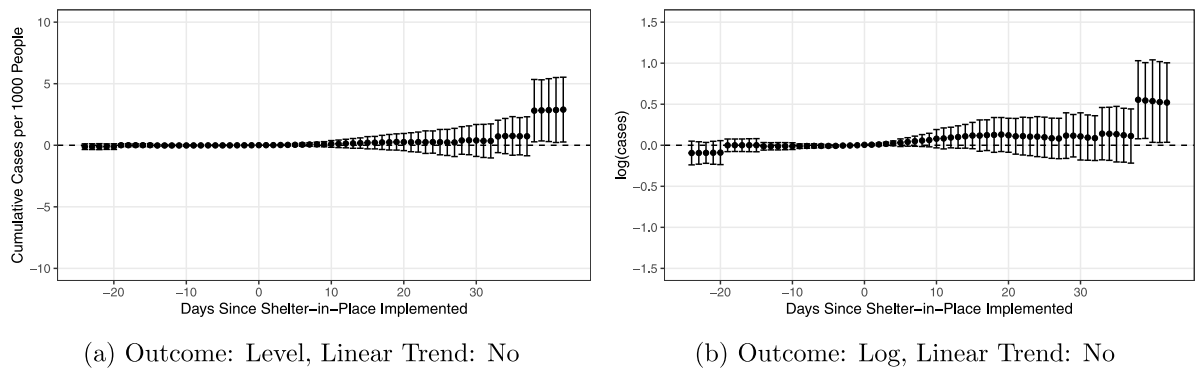
Fig. 5. Policy effects of SIPOs on Covid-19 cases using TWFE event study regressions. *Notes:* The figure contains event study type estimates of the effect of SIPOs on the number of cumulative Covid-19 cases that come from a two-way fixed effects regression. $e = 0$ corresponds to the time period when the policy was implemented. Negative values of e correspond to pre-treatment estimates of the effect of the policy (and can be thought of as pre-tests), and positive values of e correspond to estimates of the effect of the policy at different lengths of exposure to the treatment. The estimates are normalized to be equal to 0 for $e = -1$. The estimates across panels differ based on (i) whether the outcome is in levels or in logarithms and (ii) whether or not estimates include location-specific linear trends. 90% confidence intervals are provided by the vertical bars in each panel. Standard errors are clustered at the state level.

One possible concern with the previous results is related to limitations of two-way fixed effects regressions when there is variation in treatment timing and treatment effect heterogeneity (see, for example, [de Chaisemartin and D'Haultfoeuille \(2020\)](#), [Goodman-Bacon \(2021\)](#), [Sun and Abraham \(2021\)](#) and [Borusyak et al. \(2022\)](#)). There is certainly variation in treatment timing for SIPOs and heterogeneous effects (which would include things like variation in policy effects according to the time period when the policy is adopted or effects that vary with length of exposure to the policy, among others) are very likely as well. These issues are considered in [Dave et al. \(2021\)](#) and discussed extensively in [Goodman-Bacon and Marcus \(2020\)](#). In order to address these possible concerns, [Fig. 6](#) provides estimates using “heterogeneity robust” DID estimation strategies from [Callaway and Sant’Anna \(2021\)](#) and [Gardner \(2022\)](#). The setup for these results mirrors the setup from [Fig. 5](#) as the panels differ based on whether the outcome is in levels or logarithms and by whether or not the results include location-specific linear trends. The results in panels (a) and (b) that do not include a linear time trend are broadly similar to the TWFE event study regressions from above but still suggest that SIPOs *increased* Covid-19 cases. On the other hand, the results that include location-specific linear trends are notably different from the results in [Fig. 5](#). Using heterogeneity robust versions of DID estimation strategies changes the sign of the estimates and results in estimates that SIPOs increased Covid-19 cases. This is disconcerting as these estimation strategies provide a number of advantages relative to the TWFE event studies; however, as discussed above, it seems hard to understand how SIPOs could lead to an increase in Covid-19 cases. As before, we do not find large or any statistically significant estimates of Covid-19 policies in pre-treatment periods.¹²

Next, we turn to estimates using the approach based on unconfoundedness. As discussed above, these estimates require estimating a model for treatment participation and an outcome regression model. For both of these models, we include a cubic polynomial in the current number of cases per 1000 people in a state (we define the number of current cases as the change in cumulative cases over the previous seven days); we do not include the number of susceptible individuals

¹² Depending on the specification, there are either 23 or 24 pre-treatment estimates in each panel of [Fig. 6](#). Across all four panels, none of these estimates are statistically different from 0 at the 5% level.

Panel (i): Callaway and Sant'Anna (2021)



Panel (ii): Gardner (2022)

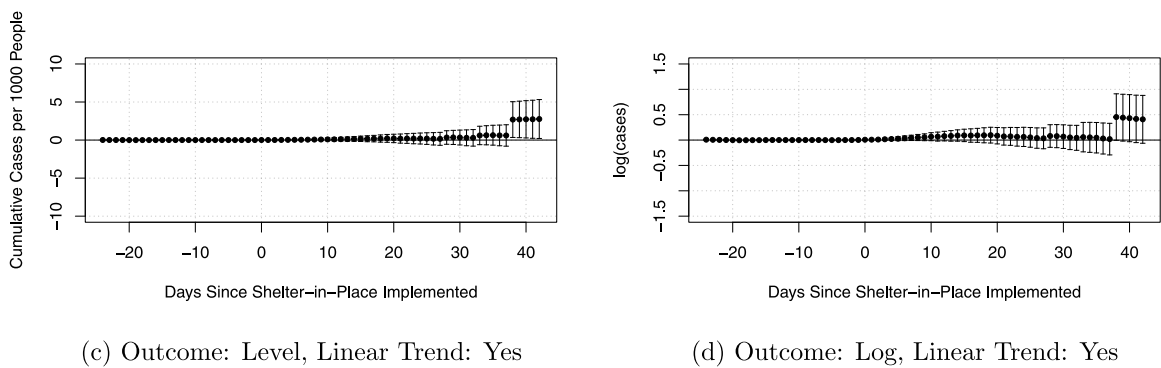


Fig. 6. Policy effects of SIPOs on Covid-19 cases using treatment effect heterogeneity-robust estimation strategies. *Notes:* The figure contains event study type estimates of the effect of SIPOs on the number of cumulative Covid-19 cases that come from several variations of heterogeneity-robust DID identification strategies (as discussed in the text). $e = 0$ corresponds to the time period when the policy was implemented. Negative values of e correspond to pre-treatment estimates of the effect of the policy (and can be thought of as pre-tests), and positive values of e correspond to estimates of the effect of the policy at different lengths of exposure to the treatment. The estimates across panels differ based on (i) whether the outcome is in levels or in logarithms and (ii) whether or not estimates include location-specific linear trends. In Panels (a) and (b) the pre-treatment estimates have the same interpretation as the TWFE estimates in Fig. 5; on the other hand, the pre-treatment estimates in Panels (c) and (d), have a placebo interpretation (i.e., they provide the estimate of the “effect of the treatment as if it had been implemented in that particular period”). Thus, the confidence intervals are systematically narrower in these panels (see Brown and Butts (2022) for additional details along these lines). 90% confidence intervals are provided by the vertical bars in each panel. Standard errors are clustered at the state level.

as this is very close to the full population in all states during the period early in the pandemic that we consider; we additionally include a dummy variable for region of the country so that states are compared to other states in the same region as well as the logarithm of the state's population; finally, we include the 7-day lag of the cumulative cases, the 7-day lag of the number of tests run per 1000 people in the state and the change in tests from the pre-treatment period to the current period in order to control for the possibility that some states were detecting Covid-19 cases better than others. When we estimate the propensity score, we find evidence that the overlap condition is violated indicating that there are a substantial number of states that do not have reasonable comparisons among never-treated and late-treated states. From this step, we drop 14 states from our analysis.¹³ The omitted states include states such as New York and California which had large numbers of early Covid-19 cases and were early adopters of SIPOs. It is, therefore, hard to find reasonable comparisons for these states. Another large number of states from the South are dropped due to the timing of their policies being very similar which makes it challenging to find reasonable comparison states when region is included as a conditioning variable.

These results are provided in Fig. 7. Using the unconfoundedness approach, we estimate relatively small and statistically insignificant effects of SIPOs.

¹³ The states that we drop are Alabama, Arizona, California, Florida, Georgia, Mississippi, Missouri, New York, North Carolina, Pennsylvania, South Carolina, Texas, Virginia, and Washington.

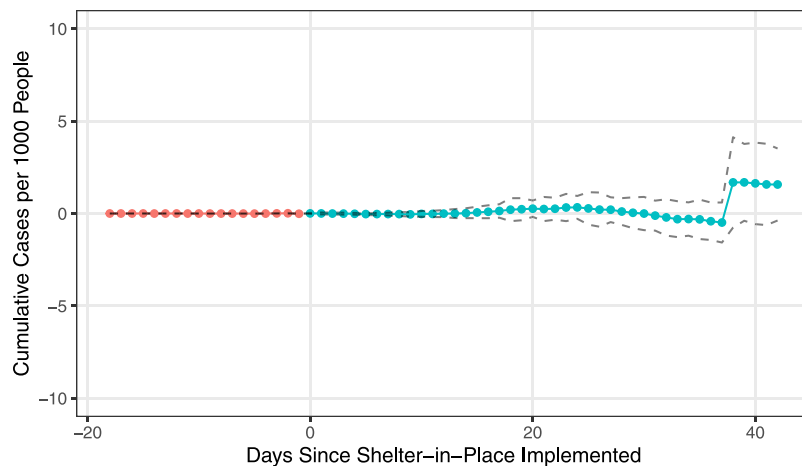


Fig. 7. Policy effects of SIPOs on Covid-19 cases under unconfoundedness. *Notes:* The figure contains event study type estimates of the effect of SIPOs on the number of cumulative Covid-19 cases uses the unconfoundedness approach discussed in the paper. $e = 0$ corresponds to the time period when the policy was implemented. Negative values of e correspond to pre-treatment estimates of the effect of the policy (and can be thought of as pre-tests), and positive values of e correspond to estimates of the effect of the policy at different lengths of exposure to the treatment. 90% confidence intervals are provided by the dashed lines. Standard errors are clustered at the state level.

7.2 County-level results

There are several complications with any state-level analysis that are worth noting. First, there are several other Covid-19 policies that were being implemented in a number of states around the same time. Several recent papers have pointed out challenges with “controlling for” other policies in linear models like the one in Eq. (16) (e.g., [de Chaisemartin and D’Haultfœuille \(2023\)](#) and [Goldsmith-Pinkham et al. \(2021\)](#)) especially in a framework (like the current one) with treatment effect heterogeneity. Moreover, using state-level data, the number of observations is very small, and, relatedly, it is challenging to find treated and untreated states that are similar enough to each other to reasonably interpret differences in outcomes as being due to the policy.¹⁴ On the other hand, at the county-level, there are often a large number of counties that are located in neighboring states and have similar characteristics, including both pre-treatment pandemic-related characteristics as well as population, median income, or demographic characteristics.¹⁵

Our strategy for the remainder of this section is to use county-level data and compare counties in states that implemented a SIPO to counties in states that did not implement a SIPO while (i) having a similar mix of other policies (both in terms of which types of policies were implemented and their timing) and (ii) making tight comparisons between counties with similar characteristics. We focus on Arkansas and Iowa which were two states that did not implement a SIPO; we compare them, in turn, to their surrounding states that implemented SIPOs.¹⁶ Among their surrounding states, Nebraska, Oklahoma, and South Dakota also did not implement a SIPO. This provides an opportunity to estimate placebo policy effects; that is, to use different estimation strategies in a case where we know that the “policy effects” should be equal to 0, which, therefore, provides a way to assess the performance of different estimation strategies.

Data

As for the state-level data, we obtain the number of county-level Covid-19 cases from the CDC COVID Data Tracker and 2019 county-level population from the Census Bureau. Some of our results below use the number of Covid-19 tests in a particular county which also comes from the CDC COVID Data Tracker.

We initially include all states that border either Arkansas or Iowa with two exceptions. We exclude Texas which shares a small border with Arkansas, and we include Kansas which “almost” borders both Arkansas and Iowa. Thus, we start with county-level data for thirteen states; these are listed in [Table 4](#).

¹⁴ To give a simple specific example, population density is likely to be an important determinant of Covid-19 transmission rates. One of the largest states that did not implement a SIPO was Arkansas, and a natural strategy is to try to compare Covid-19 cases in Arkansas to, say, Missouri which implemented a SIPO on April 6. At the state-level, though, the population density of Arkansas and Missouri is much different. Missouri’s population is double Arkansas’s population, and it has two major cities, Kansas City and St. Louis, while Arkansas does not have a major city. This suggests that, at the state-level, Arkansas may not be very useful for delivering counterfactual outcomes for Missouri had it not implemented the policy.

¹⁵ There are some potential drawbacks to using county-level data. One issue is that, by restricting comparisons to counties with similar characteristics, it potentially changes the target parameter. There are other issues related to inference such as spatial correlations and clustering standard errors at the county-level rather than the state-level as well. We discuss these issues in more detail in the Supplementary Appendix.

¹⁶ [Haynes et al. \(2022\)](#) similarly make comparisons between counties in states that implemented a SIPO to counties in states that did not. Compared to our approach, they use TWFE event study regressions and find mixed results with estimates sometimes indicating reductions in Covid-19 cases, sometimes indicating no effect, and sometimes indicating increases in Covid-19 cases.

Table 4

Timing of Covid-related policies across states.

State		Emergency Dec.	Schools closed	Shelter-in-place	Business closure	Gathering Rest.
AR	Start	March 11	March 16	–	–	March 27
	End	–	–	–	–	–
IL	Start	March 9	March 17	March 21	March 21	March 13
	End	–	–	May 30	–	May 30
KS	Start	March 12	March 17	March 30	–	March 17
	End	–	–	May 4	–	–
LA	Start	March 11	March 13	March 23	–	March 13
	End	–	–	May 15	–	May 15
IA	Start	March 9	April 3	–	–	March 17
	End	–	–	–	–	–
MN	Start	March 13	March 18	March 28	–	March 28
	End	–	–	May 18	–	–
MO	Start	March 13	March 23	April 6	April 6	March 23
	End	–	–	May 4	–	–
MS	Start	March 14	March 19	April 3	April 3	March 24
	End	–	–	April 27	–	–
NE	Start	March 13	April 1	–	–	April 3
	End	–	May 31	–	–	–
OK	Start	March 15	March 17	–	April 1	March 24
	End	May 11	–	–	–	May 20
SD	Start	March 13	March 16	–	–	April 6
	End	–	April 28	–	–	May 31
TN	Start	March 12	March 16	April 2	April 1	March 23
	End	–	–	May 1	May 30	May 30
WI	Start	March 12	March 18	March 25	March 25	March 17
	End	May 11	–	May 26	May 26	–

Notes: The table provides the timing of the main state-level policies that were implemented during the first few months of the pandemic: emergency declarations, school closures, SIPOs, business closures, and gathering restrictions. The data comes from Fullman et al. (2021) and uses their classification system of policies. Policies that were not implemented in a particular state are denoted with a “–” in the table. The end date for policies is provided for policies that were ended before the end of May; policies that did not end by May 31 are listed as “–” in the table. Some policies are somewhat ambiguously defined; this is particularly true for business closures and gathering restrictions for which there is a substantial amount of variation across states in terms of which types of businesses and gatherings were restricted. For some policies, Fullman et al. (2021) records multiple instances of the same policy. In these cases, the table provides the first state-wide, mandated version of the policy, and the end date is the period when there was no longer a state-wide, mandated policy (in some cases, particularly for business closures and gathering restrictions, this can involve overlapping policies with variation in the intensity of the restrictions). In the Supplementary Appendix, we provide some additional details regarding some particular policies.

A major concern for our application is the timing of other Covid-19 related policies across states. Table 4 provides the date when each state declared an emergency, closed schools, implemented a SIPO, closed non-essential businesses, or imposed gathering restrictions. Overall, Arkansas tended to implement other policies with very similar timing as all of its surrounding states especially with respect to the timing of the emergency declaration and school closures. There is more variation with respect to non-essential business closures and gathering restrictions though we note that these policies are less clearly defined than the other main Covid-related policies. Our interpretation is that it is reasonable to consider the mix and timing of other policies as being quite similar to all of its neighboring states. The timing of Iowa's policies are somewhat more different from its surrounding states. The primary difference is that Iowa closed its schools somewhat later than surrounding states with the exception of Nebraska. This arguably suggests that the results below that use Iowa as the comparison state are somewhat less credible (with a viable alternative explanation that differences in school closure policy could be at least partially driving the results).

A second major concern is being able to find counties with similar pre-SIPO pandemic related characteristics. For each pair of states, we use the same model for the outcome regression and propensity score as for the state-level data that includes a cubic polynomial in the current number of cases, the 7-day lag of cumulative cases, the 7-day lag and change in the number of Covid-19 tests in the county, and the logarithm of county population. In order to enforce the overlap condition between treated and untreated states, we drop treated counties with an estimated propensity score greater than 0.95.¹⁷ With the remaining set of counties, including both treated and untreated counties, we compute the propensity score weights from Theorem 2 using the same set of covariates as mentioned above. Given these weights, we compute the effective sample size for untreated states, which is defined as $\sum_{i \in \mathcal{U}} w_i^2 / \left(\sum_{i \in \mathcal{U}} w_i^2 \right)$ where \mathcal{U} denotes the set of untreated

¹⁷ In practice, for our main results, we do not drop many counties from this procedure. When we compare Missouri to Arkansas, we do not drop any counties using this criteria. When we compare Minnesota to Iowa, we drop Ramsey County which is the county where St. Paul is located. As discussed above, dropping counties this way does change the target parameter from being the ATT to being an ATT-type parameter that is local to the region of common support. And, for example, our approach tends to drop urban counties relative to suburban or rural counties; if SIPOs tended to reduce Covid-19 cases more in urban counties than in other counties, then our estimation procedure would tend to result in smaller in magnitude estimates of policy effects.

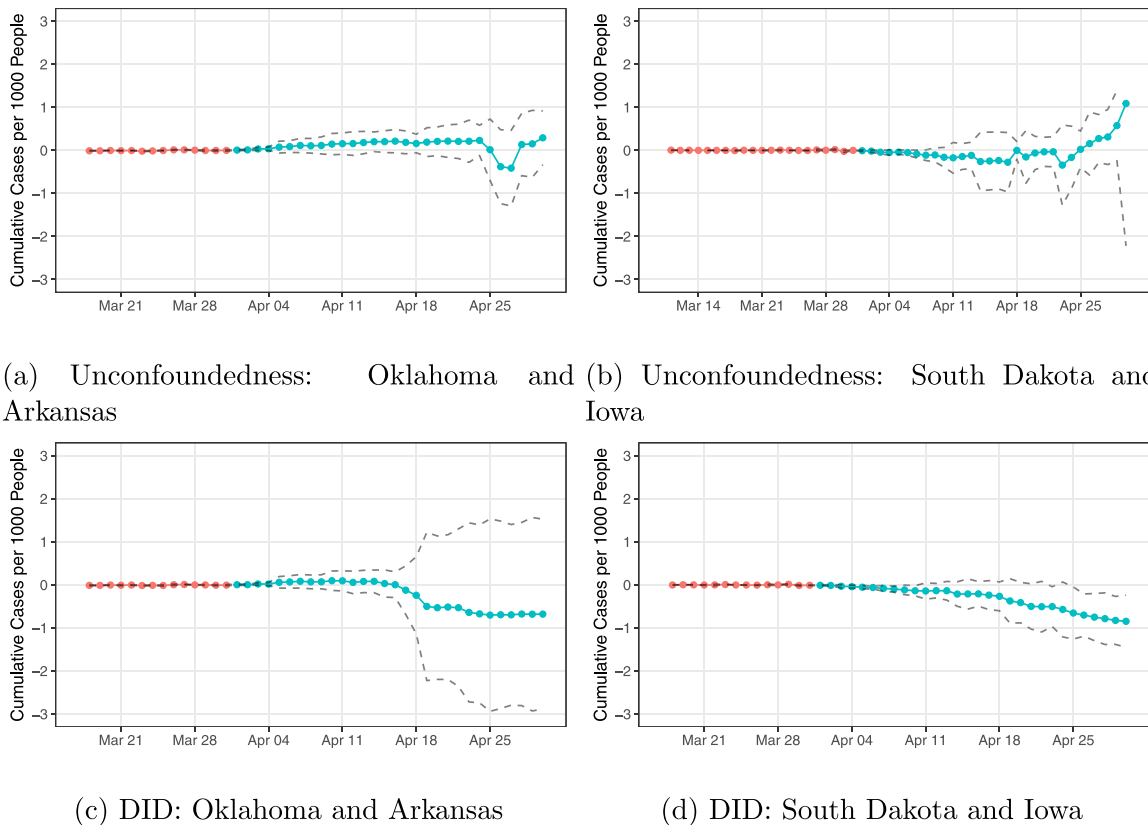


Fig. 8. Placebo policy effects. *Notes:* The figure provides estimates of placebo SIPOs (i.e., policies that were not actually implemented) for Oklahoma and South Dakota using the unconfoundedness and DID approaches discussed in the main text. The placebo policy date is set to be April 1 for both Oklahoma and South Dakota. Estimates before the placebo policy date are in red while estimates in periods after the placebo policy date are provided in blue. The dashed lines provide 90% confidence intervals. Standard errors are clustered at the county level. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

county (see, for example, [Kish \(1965\)](#), [Chattopadhyay and Zubizarreta \(2022\)](#) and [Shook-Sa and Hudgens \(2022\)](#)). Then, we exclude combinations of states where the effective sample size is less than or equal to 25 counties for either the treated or untreated state. Following this procedure, we end up with four combinations of states that satisfy this criteria. Two of these, (i) Oklahoma relative to Arkansas and (ii) South Dakota relative to Iowa, arise for states that did not implement a SIPO. We estimate placebo policy effects for these states. The other two combinations of states are (iii) Missouri relative to Arkansas and (iv) Minnesota relative to Iowa. These states provide a chance to estimate effects of policies that were actually implemented.

Results

We start by providing placebo policy effect estimates for Oklahoma and South Dakota – neither of which implemented a SIPO. The reason for starting here is that this serves as a good way to evaluate estimation strategies; policy effects should, in principle, be equal to 0 in all time periods. In order to facilitate comparisons to other results, for states that did not implement a SIPO, we use a “placebo policy” date of April 1. These results are provided in [Fig. 8](#). Panels (a) and (b) provide estimates using the unconfoundedness approach. These estimates are generally small and not statistically significantly different from zero. Estimates using a difference-in-differences approach are provided in panels (c) and (d). For Oklahoma, none of the post-policy estimates are statistically different from zero though the standard errors are notably larger using DID than for the unconfoundedness approach. More notably, using DID, we spuriously estimate that the placebo policy in South Dakota reduced the number of Covid-19 cases in South Dakota. These results suggest that, at least for these two policies, the unconfoundedness approach performs better than DID. It is also worth mentioning that, due to the relatively similar pre-pandemic characteristics of Oklahoma to Arkansas and South Dakota to Iowa, this is still a relatively favorable setting for DID estimation strategies.

Next, we move to estimating effects of policies that were actually implemented in Missouri (using counties from Arkansas as the comparison) and in Minnesota (using counties from Iowa as the comparison). The results under unconfoundedness are provided in [Fig. 9](#). For Missouri, we do not find evidence of policy effects on Covid-19 cases. On the

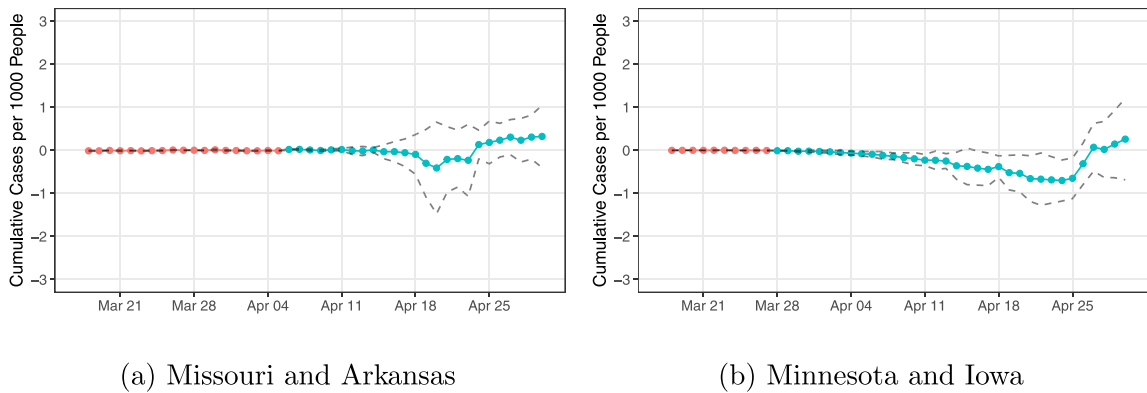


Fig. 9. County-level estimates under unconfoundedness. *Notes:* The figure provides estimates of SIPOs effects on Covid-19 cases for Missouri and Minnesota using the unconfoundedness approach discussed in the main text. Pre-treatment estimates are reported in red while post-treatment estimates are reported in blue. The pre-treatment estimates use the immediately preceding day as the base period while all the post-treatment periods use the period immediately before the treatment as the base period. The dashed lines provide 90% confidence intervals. Standard errors are clustered at the county level. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

other hand, for Minnesota, we estimate that its SIPO decreased Covid-19 cases at least in some periods. For example, on April 25, we estimate that Minnesota's SIPO reduced Covid-19 cases by about 0.46 cases per 1000 people in the state. In results provided in the Supplementary Appendix, using DID, we estimate a large reduction in Covid-19 cases in Missouri due to the policy though the standard errors are much larger than for the results based on unconfoundedness and not statistically different from 0. For Minnesota, the DID estimates are quite similar to the unconfoundedness results presented here.

In the Supplementary Appendix, we provide additional estimates for all combinations of states discussed in this section. Of the remaining 8 combinations of states that (i) implemented a policy and (ii) exhibit similar pre-treatment paths of Covid-19 cases, using the same DID estimation strategy used in this section, two of the estimates are positive and statistically significant, two of the estimates are negative and statistically significant, and four of the estimates are not statistically different from zero. Thus, like the case with state-level data above, using DID can lead to hard-to-explain positive estimates of SIPOs on Covid-19 cases and, more generally, estimates that seem inconsistent with each other. Arguably, these differences could be explained by heterogeneous effects of policies implemented in different states, though, in our view, the large magnitude of the differences in estimated policy effects cuts against heterogeneous policy effects being a full explanation of the differences. Instead, a better explanation seems to be that these differences are driven by pre-treatment differences in the state of the pandemic between counties in states that implemented SIPOs relative to counties in neighboring states that did not.

Finally, we consider the effect of SIPOs on travel. We focus on the percentage change in retail and recreation travel from a pre-Covid baseline. In the main text, we focus on the effect of the policy on travel in Missouri using Arkansas as the comparison state. These results are available in Fig. 10. The results in Panel (a) come from a standard DID approach that implicitly imposes that Covid-19 cases do not directly affect travel; the results in Panel (b) come from the regression DID approach that includes current cases as a covariate but not accounting for the possibility that SIPOs could have affected the number of cases directly; and the results in Panel (c) use the adjusted regression DID approach proposed in the current paper that allows for the policy to have had an effect on Covid-19 cases.

In this case, the estimates are more broadly similar than they were for cumulative Covid-19 cases. Standard DID estimates indicate a relatively small but persistent negative effect of SIPOs on retail and recreation travel. Using standard DID, the overall estimate of the effect of SIPOs across post-treatment time periods is -1.30 (p-value: 0.16); i.e., across the first several weeks of the SIPO, the policy reduced travel by about 1.3 percent relative to what travel would have been if the policy had not been implemented. The point estimates from regression DID (these estimates just include observed current cases as a covariate) are somewhat larger in magnitude; in this case, the estimated overall effect of SIPOs on travel is -1.63 (p-value: 0.08). Finally, the adjusted regression DID overall estimated effect of SIPOs is -1.25 (p-value: 0.19) when we allow for the policy to have had an effect on cases. In the Supplementary Appendix, we provide analogous results using the state-level data from earlier in this section. Those results are broadly similar to the ones presented here though we estimate somewhat larger-in-magnitude effects (around a 3 percent reduction in travel due to SIPOs) that do not vary much by estimation strategy and are somewhat more precisely estimated than the results presented here.

7.3 Discussion of results

Our results, especially those about the effect of SIPOs on the number of Covid-19 cases, are substantially different from existing estimates, and it is worth making a few additional comments. First, our results on the effect of SIPOs on travel

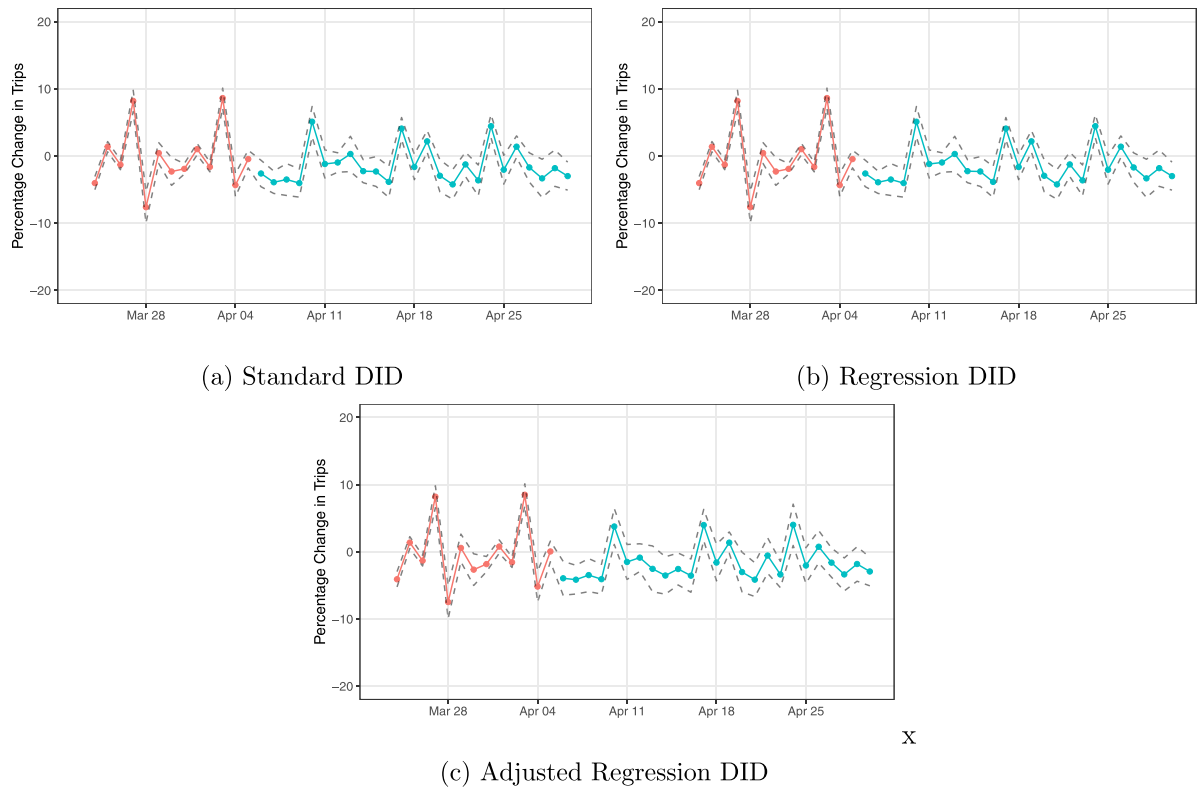


Fig. 10. Estimates of SIPO orders on travel. *Notes:* The figure contains event study type estimates of the effect of SIPOs on the percentage change in retail and recreation travel. $e = 0$ corresponds to the time period when the policy was implemented. Negative values of e correspond to pre-treatment estimates of the effect of the policy and can be thought of as pre-tests, and positive values of e correspond to estimates of the effect of the policy at different lengths of exposure to the treatment. Panel (a) provides estimates using standard DID (without accounting for cases), Panel (b) provides regression DID estimates (accounting for cases but not that the policy may have a direct effect on cases), and Panel (c) provides adjusted regression DID estimates (accounting for cases and allowing for the policy to have had an effect on cases as is proposed in the text). The dashed line provides 90% confidence intervals. Standard errors are clustered at the county level.

are more similar to existing estimates (e.g., [Goolsbee and Syverson \(2021\)](#)). Since the primary channel through which SIPOs would likely reduce Covid-19 cases is through reducing travel/contact with other individuals, it seems reasonable to simultaneously estimate relatively small effects of SIPOs on travel coinciding with small effects of SIPOs on the number of Covid-19 cases; but harder to rationalize SIPOs strongly decreasing Covid-19 cases while having only a small effect on travel.

That being said, we hesitate to interpret our results as providing strong evidence that SIPOs did not reduce Covid-19 cases. For one thing, the interpretation of our treatment effect parameters is somewhat subtle. Untreated potential outcomes here do not correspond to particular locations not reacting at all to Covid-19 but rather to outcomes that would have occurred if a state had not implemented the policy (but other things about the state remained the same). This can be seen to be clearly relevant from the summary statistics in [Table 3](#) where all states (not just those who implemented the policy) were experiencing massive decreases in travel over the period that we consider. Importantly, this indicates that our results are not at all saying that staying at home did not have an effect on Covid-19 cases. Second, the statistical power of all of the estimation procedures considered above is relatively low (this is true both for the unconfoundedness approach and the DID approaches) and leads to generally wide confidence intervals on estimated policy effects. And even though a number of our estimates are not statistically significant, many of them are large enough to be compatible with a wide variety of possible effects of SIPOs on Covid-19 cases.¹⁸ Instead, we interpret the results from the application as indicating that these sorts of policies are likely to be very challenging to precisely evaluate due to both data limitations as

¹⁸ To give an example, in our results for Missouri, which is a state for which we do not estimate a statistically significant policy effects on Covid-19 cases in post-treatment periods (see [Fig. 9](#)), the lower end of a 90% confidence interval for our estimated effect of the policy on Covid-19 cases 21 days after the policy was implemented is that it reduced Covid-19 cases by about 0.13 per 1000 people. If we divide this by the estimated number of Covid-19 cases that treated locations would have experienced in the same period if they had not implemented the policy (i.e., $\mathbb{E}[C_t(0)|D = 1]$ in [Eq. \(6\)](#) which is identified and can be recovered in our setup), we would estimate that SIPOs decreased Covid-19 cases by 30%. In other words, our estimates do not necessarily rule out the possibility that SIPOs may have had quite large effects on Covid-19 cases.

well as trying to deal with a highly nonlinear outcome and a policy that was adopted at different times by states whose exposure to the pandemic also varied widely and was correlated with the timing of the policy being adopted.

8 Conclusion

In this paper, we have considered several different policy evaluation strategies and how compatible they are with a leading epidemic model. For identifying the direct effects of policies on the number of Covid-19 cases, our results suggest that strategies based on unconfoundedness type assumptions are likely to perform better than difference-in-differences type strategies due to the highly nonlinear nature of the spread of Covid-19.

Our second main set of results were about evaluating the effects of policies on other economic outcomes when (i) the policy can affect the number of Covid-19 cases and (ii) the number of Covid-19 cases can have a direct effect on the economic outcome of interest. For this case, we also showed that two of the most common ways to evaluate these policies (difference-in-differences directly or including the number of cases as a covariate in a DID setup) do not generally deliver an average effect of the policy. We proposed an alternative estimator that is valid in this case.

We applied our approach to study the effects of shelter-in-place orders on Covid-19 cases and travel early in the pandemic. We showed that our theoretical arguments were indeed relevant in this context and led to notably different estimates (particularly for the number of Covid-19 cases) relative to the most common approaches used in applications.

There remain a number of interesting possible extensions to this work, and we conclude by mentioning two of them. First, evaluating the effects of various policies (especially early in the pandemic) is complicated by limited and nonrandom Covid-19 testing during the early part of the pandemic (see, for example, Callaway and Li (2023) and Manski and Molinari (2020)), and it would be interesting to extend our results along these dimensions. Second, another common policy evaluation approach in the context of Covid-19 related policies is the synthetic control method (examples include Cho (2020), Dave et al. (2020a), Friedson et al. (2021) and Mitze et al. (2020)). It appears that, most often, a researcher's decision between difference-in-differences or synthetic controls is driven by whether the number of treated locations is large or small. However, it is less clear under what conditions synthetic control approaches are compatible with the sorts of epidemiological models that we considered in the current paper.

Appendix A. More details on stochastic SIRD models

A.1. Stochastic SIRD model for untreated potential outcomes

In this section, we write down a stochastic SIRD model along the lines of Eqs. (1)–(5) but for untreated potential outcomes and written in an error form.

Stochastic SIRD Model for Untreated Potential Outcome. For all $t = 2, \dots, T$

$$I_{lt}(0) = I_{lt-1}(0) + \left(\beta \frac{I_{lt-1}(0)}{N_l} S_{lt-1}(0) + u_{lt}^I \right) - \Delta R_{lt}(0) - \Delta \delta_{lt}(0) \quad (\text{A.1})$$

$$R_{lt}(0) = R_{lt-1}(0) + \lambda I_{lt-1}(0) + u_{lt}^R \quad (\text{A.2})$$

$$\delta_{lt}(0) = \delta_{lt-1}(0) + \gamma I_{lt-1}(0) + u_{lt}^\delta \quad (\text{A.3})$$

$$N_l = S_{lt}(0) + I_{lt}(0) + R_{lt}(0) + \delta_{lt}(0) \quad (\text{A.4})$$

$$C_{lt}(0) = N_l - S_{lt}(0) \quad (\text{A.5})$$

where we omit an equation for $S_{lt}(0)$, the number of susceptible in location l in time period t , because it is fully determined by $I_{lt}(0)$, $R_{lt}(0)$, $\delta_{lt}(0)$, and N_l . Next, define $u_{lt} := (u_{lt}^I, u_{lt}^R, u_{lt}^\delta)'$. We also make the following assumptions about u_{lt}

$$\mathbb{E}[u_t | \mathcal{F}_{t-1}(0), \dots, \mathcal{F}_1(0), D = d] = 0 \text{ for all } t = 2, \dots, T \quad (\text{A.6})$$

and that

$$u_t \perp (D, \mathcal{F}_{t-2}(0), \dots, \mathcal{F}_1(0)) | \mathcal{F}_{t-1}(0) \text{ for all } t = 2, \dots, T \quad (\text{A.7})$$

Together with Eqs. (A.1)–(A.5), Eq. (A.6) completes the stochastic SIRD model for untreated potential outcomes in an analogous way to the model presented in the main text. Note that, together, these conditions rule out that u_t are serially correlated. Additionally, Eqs. (A.6) and (A.7) are related to the Markov property of stochastic SIRD models discussed in the main text in that the distribution of error terms can depend on the state of the pandemic in the immediately preceding period (this allows, for example, the variance of the error terms to depend on the state of the pandemic in the previous period), but their distribution does not depend on further lags of the state of the pandemic. These conditions also have the flavor of sequential exogeneity assumptions in the panel data econometrics literature where the distribution of error terms does not depend on lags of the state of the pandemic (besides the first one), but current shocks can and do impact future pandemic states. Both Eqs. (A.6) and (A.7) also involve error terms being independent

Table 5
Simulation parameters.

Parameter	Notation	Value
Infection rate	β	0.08
Post policy β	β_{pol}	0.08
Recovery rate	λ	0.04
Death rate	γ	0.003
Time periods	T	400
Population size	N	1000
Treatment probability	p	0.5
Initial case period (treated)		Poisson($\lambda = 40$)
Initial case period (untreated)		Poisson($\lambda = 80$)
Initial # cases		10
# locations		250
Policy start period		150

of treatment status conditional on pre-treatment status. This is arguably the key condition here and indicates that, given the state of the pandemic in period $t - 1$, the distribution of shocks for treated locations is the same as for untreated locations. Relative to typical implementations of stochastic SIRD models, these are mild conditions. For example, it is typical to impose that, $\left(\beta \frac{I_{t-1}(0)}{N_l} S_{lt-1}(0) + u_{lt}^I\right)$, which is the number of new infections, follows a Poisson distribution with parameter $\beta \frac{I_{t-1}(0)}{N_l} S_{lt-1}(0)$, and that the number of new recoveries $(\lambda I_{lt-1}(0) + u_{lt}^R)$, new deaths $(\gamma I_{lt-1}(0) + u_{lt}^D)$, and continued infections follows a multinomial distribution with parameters $(\lambda I_{lt-1}(0), \gamma I_{lt-1}(0), (1 - \lambda - \gamma) I_{lt-1}(0))$. This set up is compatible with the Stochastic SIRD Model for Untreated Potential Outcomes discussed above.

A.2. Simulation details

Table 5 provides the values of the parameters that we use in the simulations. The most important things to notice are that (i) the policy has no effect on Covid-19 infections rates (i.e., $\beta = \beta_{pol}$) and (ii) the distribution of the timing of the first case is different across locations that participate in the treatment and those that do not. This is roughly analogous to the idea that locations that were exposed to Covid-19 earlier tended to implement policies earlier.

Appendix B. Proofs

Proof of Theorem 1

Proof. For the first part, notice that

$$\begin{aligned}
 \mathbb{E}[\Delta^{(t^*-1,t)} C_t(0) | D = d] &= \mathbb{E}[C_t(0) - C_{t^*-1}(0) | D = d] \\
 &= \sum_{s=t^*}^t \mathbb{E}[\Delta C_s(0) | D = d] \\
 &= \sum_{s=t^*}^t \mathbb{E}[\mathbb{E}[\Delta C_s(0) | \mathcal{F}_{t^*-1}, D = d] | D = d]
 \end{aligned} \tag{B.1}$$

where the first equality holds by the definition of $\Delta^{(t^*-1,t)} C_t(0)$, the second equality holds by adding and subtracting $\mathbb{E}[C_s(0) | D = d]$ for all $s = t^*, \dots, (t - 1)$, and the last equality holds by the law of iterated expectations.

For the second part,

$$\begin{aligned}
 ATT_t^C &= \mathbb{E}[C_t(1) - C_t(0) | D = 1] \\
 &= \mathbb{E}[C_t(1) - C_{t^*-1}(0) | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} C_t(0) | D = 1] \\
 &= \mathbb{E}[C_t(1) - C_{t^*-1}(0) | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} C_t(0) | D = 0] \\
 &\quad + \left(\mathbb{E}[\Delta^{(t^*-1,t)} C_t(0) | D = 0] - \mathbb{E}[\Delta^{(t^*-1,t)} C_t(0) | D = 1] \right) \\
 &= \mathbb{E}[\Delta^{(t^*-1,t)} C_t | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} C_t | D = 0] \\
 &\quad + \left(\sum_{s=t^*}^t \mathbb{E}[\mathbb{E}[\Delta C_s(0) | \mathcal{F}_{t^*-1}, D = 0] | D = 0] - \sum_{s=t^*}^t \mathbb{E}[\mathbb{E}[\Delta C_s(0) | \mathcal{F}_{t^*-1}, D = 1] | D = 1] \right)
 \end{aligned}$$

where the first equality comes from the definition of ATT_t^C , the second equality holds by adding and subtracting $\mathbb{E}[C_{t^*-1}(0)|D = 1]$ (which is the average number of Covid-19 cases across treated locations in the pre-treatment period), the third equality holds by adding and subtracting $\mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = 0]$ (which is the average path of Covid-19 cases from period $t^* - 1$ to t among untreated locations), and the fourth equality holds by rewriting potential outcomes in terms of their observed counterparts and by Eq. (B.1). This implies the second result in the theorem. \square

Proof of Proposition 1

To conserve on notation, define the function $r(\mathcal{F}_t(0), u_t)$ to be the vector-valued function where the first argument takes in the state of the pandemic in some period t , $\mathcal{F}_t(0)$, and vector of error terms u_t and returns $\mathcal{F}_{t+1}(0)$ according to the Stochastic SIRD Model for Untreated Potential Outcomes. Then, for some post-treatment time period $t \geq t^*$, we have that

$$\mathcal{F}_{lt} = r(\mathcal{F}_{lt-1}, u_{lt}) \quad (\text{B.2})$$

Before proving Proposition 1, we provide two helpful intermediate results.

Lemma 1. *In the Stochastic SIRD Model for Untreated Potential Outcomes and under Assumption 1, for any $t \geq t^*$, for a function h , $d \in \{0, 1\}$, and for some $k > 0$ such that $t - k \geq 1$,*

$$\mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), \dots, \mathcal{F}_{t-k}(0), D = d] = \mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), D = d]$$

Proof. Notice that

$$\begin{aligned} \mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), \dots, \mathcal{F}_{t-k}(0), D = d] &= \mathbb{E}[h(r(\mathcal{F}_{t-1}(0), u_t))|\mathcal{F}_{t-1}(0), \dots, \mathcal{F}_{t-k}, D = d] \\ &= \mathbb{E}[h(r(\mathcal{F}_{t-1}(0), u_t))|\mathcal{F}_{t-1}(0), D = d] \\ &= \mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), D = d] \end{aligned}$$

where the first and last equalities holds by the definition of r , and the second equality holds by Eq. (A.7) in the Stochastic SIRD Model for Untreated Potential Outcomes. \square

Lemma 2. *In the Stochastic SIRD Model for Untreated Potential Outcomes and under Assumption 1, for any $t \geq t^*$ and for a function h ,*

$$\mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), D = 1] = \mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), D = 0]$$

Proof. Notice that

$$\begin{aligned} \mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), D = 1] &= \mathbb{E}[h(r(\mathcal{F}_{t-1}(0), u_t))|\mathcal{F}_{t-1}(0), D = 1] \\ &= \mathbb{E}[h(r(\mathcal{F}_{t-1}(0), u_t))|\mathcal{F}_{t-1}(0), D = 0] \\ &= \mathbb{E}[h(\mathcal{F}_t(0))|\mathcal{F}_{t-1}(0), D = 0] \end{aligned}$$

where the first and last equalities hold by the definition of r and the second equality holds by Eq. (A.7) in the Stochastic SIRD Model for Untreated Potential Outcomes. \square

Proof of Proposition 1. For the proof, we provide the more general result that $\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t^*-1}(0), D = 1] = \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t^*-1}(0), D = 0]$; i.e., we prove unconfoundedness holds for all the variables in the epidemic model which includes cumulative cases, but also includes the number of susceptible, the number of currently infected, and the number of recovered and deaths.

For $k \geq 1$, we recursively define

$$g_d^{(k)}(\mathcal{F}_{s-k}(0)) := \begin{cases} \mathbb{E}[g_d^{(k-1)}(\mathcal{F}_{s-k+1}(0))|\mathcal{F}_{s-k}(0), D = d] & \text{if } k > 1 \\ \mathbb{E}[\mathcal{F}_s(0)|\mathcal{F}_{s-1}(0), D = d] & \text{if } k = 1 \end{cases}$$

Using this notation, notice that we can write

$$\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), D = 1] = g_1^{(1)}(\mathcal{F}_{t-1}(0)) \quad (\text{B.3})$$

and that

$$\begin{aligned} \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-2}(0), D = 1] &= \mathbb{E}\left[\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), \mathcal{F}_{t-2}(0), D = 1] \middle| \mathcal{F}_{t-2}(0), D = 1\right] \\ &= \mathbb{E}\left[\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), D = 1] \middle| \mathcal{F}_{t-2}(0), D = 1\right] \\ &= \mathbb{E}\left[g_1^{(1)}(\mathcal{F}_{t-1}(0)) \middle| \mathcal{F}_{t-2}(0), D = 1\right] \end{aligned} \quad (\text{B.4})$$

$$= g_1^{(2)}(\mathcal{F}_{t-2}(0)) \quad (\text{B.5})$$

where the first equality holds by the law of iterated expectations, the second equality by [Lemma 1](#), the third equality by the definition of $g_1^{(1)}$, and the last equality by the definition of $g_1^{(2)}$. Similarly, notice that

$$\begin{aligned}
 \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-3}(0), D = 1] &= \mathbb{E}\left[\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), \mathcal{F}_{t-2}(0), \mathcal{F}_{t-3}(0), D = 1] \middle| \mathcal{F}_{t-3}(0), D = 1\right] \\
 &= \mathbb{E}\left[\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), D = 1] \middle| \mathcal{F}_{t-3}(0), D = 1\right] \\
 &= \mathbb{E}\left[g_1^{(1)}(\mathcal{F}_{t-1}(0)) \middle| \mathcal{F}_{t-3}(0), D = 1\right] \\
 &= \mathbb{E}\left[\mathbb{E}[g_1^{(1)}(\mathcal{F}_{t-1}(0))|\mathcal{F}_{t-2}(0), \mathcal{F}_{t-3}(0), D = 1] \middle| \mathcal{F}_{t-3}(0), D = 1\right] \\
 &= \mathbb{E}\left[\mathbb{E}[g_1^{(1)}(\mathcal{F}_{t-1}(0))|\mathcal{F}_{t-2}(0), D = 1] \middle| \mathcal{F}_{t-3}(0), D = 1\right] \\
 &= \mathbb{E}\left[g_1^{(2)}(\mathcal{F}_{t-2}(0)) \middle| \mathcal{F}_{t-3}(0), D = 1\right] \\
 &= g_1^{(3)}(\mathcal{F}_{t-3}(0))
 \end{aligned}$$

where the first equality holds by the law of iterated expectations, the second equality holds by [Lemma 1](#), the third equality holds by the definition of $g_1^{(1)}$, the fourth equality holds by the law of iterated expectations, the fifth equality holds by [Lemma 1](#), the sixth equality holds by the definition of $g_1^{(2)}$, and the last equality holds by the definition of $g_1^{(3)}$. Continuing to proceed in this way, we have that

$$\begin{aligned}
 \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t^*-1}(0), D = 1] &= \mathbb{E}\left[g_1^{(t-t^*)}(\mathcal{F}_t^*(0)) \middle| \mathcal{F}_{t^*-1}(0), D = 1\right] \\
 &= g_1^{(t-t^*+1)}(\mathcal{F}_{t^*-1}(0))
 \end{aligned}$$

The same sorts of arguments can be used to show that $\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t^*-1}(0), D = 0] = g_0^{(t-t^*+1)}(\mathcal{F}_{t^*-1}(0))$.

Next, notice that [Lemma 2](#), together with [Eq. \(B.3\)](#) implies that

$$g_1^{(1)}(\mathcal{F}_{t-1}(0)) = \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), D = 1] = \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t-1}(0), D = 0] = g_0^{(1)}(\mathcal{F}_{t-1}(0)) \quad (\text{B.6})$$

Similarly,

$$\begin{aligned}
 g_1^{(2)}(\mathcal{F}_{t-2}(0)) &= \mathbb{E}\left[g_1^{(1)}(\mathcal{F}_{t-1}(0)) \middle| \mathcal{F}_{t-2}(0), D = 1\right] \\
 &= \mathbb{E}\left[g_1^{(1)}(\mathcal{F}_{t-1}(0)) \middle| \mathcal{F}_{t-2}(0), D = 0\right] \\
 &= \mathbb{E}\left[g_0^{(1)}(\mathcal{F}_{t-1}(0)) \middle| \mathcal{F}_{t-2}(0), D = 0\right] \\
 &= g_0^{(2)}(\mathcal{F}_{t-2}(0))
 \end{aligned}$$

where the first and last equalities hold by the definition of $g_d^{(2)}$, the second equality holds by [Lemma 2](#), and the third equality holds by [Eq. \(B.6\)](#).

Proceeding along these lines, it follows that

$$g_1^{(k)}(\mathcal{F}_{t-k}(0)) = g_0^{(k)}(\mathcal{F}_{t-k}(0)) \text{ for all } k = 1, \dots, (t - t^* + 1)$$

Taking $k = t - t^* + 1$ implies that $\mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t^*-1}(0), D = 1] = \mathbb{E}[\mathcal{F}_t(0)|\mathcal{F}_{t^*-1}(0), D = 0]$ which completes the proof. \square

Proof of [Theorem 2](#)

Proof. Given the result in [Proposition 1](#), the proof of this result holds under standard arguments for unconfoundedness. We provide them here for completeness. First, notice that

$$\begin{aligned}
 ATT_t^C &= \mathbb{E}[C_t(1) - C_t(0)|D = 1] \\
 &= \mathbb{E}[C_t|D = 1] - \mathbb{E}[\mathbb{E}[C_t|\mathcal{F}_{t^*-1}, D = 0]|D = 1]
 \end{aligned} \quad (\text{B.7})$$

where the first equality holds from the definition of ATT_t^C , and the second equality holds by [Proposition 1](#). This completes the first part of the proof. For the second part, notice that continuing from Eq. (B.7),

$$\begin{aligned} ATT_t^C &= \mathbb{E} \left[\frac{D}{\mathbb{E}[D]} C_t \right] - \mathbb{E} \left[\frac{\mathbb{E}[(1-D)C_t | \mathcal{F}_{t^*-1}]}{1 - p(\mathcal{F}_{t^*-1})} \middle| D = 1 \right] \\ &= \mathbb{E} \left[\frac{D}{\mathbb{E}[D]} C_t \right] - \mathbb{E} \left[\frac{p(\mathcal{F}_{t^*-1})}{\mathbb{E}[D](1 - p(\mathcal{F}_{t^*-1}))} \mathbb{E}[(1-D)C_t | \mathcal{F}_{t^*-1}] \right] \\ &= \mathbb{E} \left[\frac{D}{\mathbb{E}[D]} C_t \right] - \mathbb{E} \left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{\mathbb{E}[D](1 - p(\mathcal{F}_{t^*-1}))} C_t \right] \end{aligned} \quad (\text{B.8})$$

Further, notice that

$$\begin{aligned} \mathbb{E} \left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{\mathbb{E}[D](1 - p(\mathcal{F}_{t^*-1}))} \right] &= \mathbb{E} \left[\mathbb{E} \left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{\mathbb{E}[D](1 - p(\mathcal{F}_{t^*-1}))} \middle| \mathcal{F}_{t^*-1} \right] \right] \\ &= \frac{\mathbb{E}[p(\mathcal{F}_{t^*-1})]}{\mathbb{E}[D]} \\ &= 1 \end{aligned} \quad (\text{B.9})$$

Combining the results from Eqs. (B.8) and (B.9) implies that

$$ATT_t^C = \mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})C_t] \quad (\text{B.10})$$

To conclude the proof, notice that

$$\begin{aligned} \mathbb{E}[\omega(D, \mathcal{F}_{t^*-1}) | \mathcal{F}_{t^*-1}] &= \frac{1}{\mathbb{E}[D]} \left(\mathbb{E}[D | \mathcal{F}_{t^*-1}] - \mathbb{E} \left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{(1 - p(\mathcal{F}_{t^*-1}))} \middle| \mathcal{F}_{t^*-1} \right] \right) \\ &= 0 \end{aligned} \quad (\text{B.11})$$

which holds from the definition of ω and also uses the argument in Eq. (B.9). This implies that

$$\begin{aligned} \mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})m_{0,t}^C(\mathcal{F}_{t^*-1})] &= \mathbb{E}[m_{0,t}^C(\mathcal{F}_{t^*-1})\mathbb{E}[\omega(D, \mathcal{F}_{t^*-1}) | \mathcal{F}_{t^*-1}]] \\ &= 0 \end{aligned} \quad (\text{B.12})$$

Combining Eq. (B.8) and Eq. (B.12) implies the second part of the result. \square

Proof of [Theorem 3](#)

Proof. For the first part, notice that

$$\begin{aligned} \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] &= \mathbb{E}[I_t(0)|D=1] - \mathbb{E}[I_{t^*-1}|D=1] \\ &= \mathbb{E}[I_t(1)|D=1] - (\mathbb{E}[I_t(1)|D=1] - \mathbb{E}[I_t(0)|D=1]) - \mathbb{E}[I_{t^*-1}|D=1] \\ &= \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D=1] - ATT_t^I \\ &= \mathbb{E} \left[\frac{D}{\mathbb{E}[D]} \Delta^{(t^*-1,t)}I_t - \omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^C(\mathcal{F}_{t^*-1})) \right] \end{aligned}$$

where the first equality holds by splitting the difference, the second equality adds and subtracts $\mathbb{E}[I_t(1)|D=1]$, the third equality holds by combining terms and the definition of ATT_t^I , and the last equality holds by using the same arguments for ATT_t^I as were used for ATT_t^C in [Theorem 2](#).

For the second part, to start with, notice that

$$\begin{aligned} ATT_t^Y &= \mathbb{E}[Y_t(1) - Y_t(0)|D=1] \\ &= \mathbb{E}[Y_t(1) - Y_{t^*-1}(0)|D=1] - \mathbb{E}[Y_t(0) - Y_{t^*-1}(0)|D=1] \\ &= \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=1] \end{aligned} \quad (\text{B.13})$$

where the first equality holds by the definition of ATT_t^Y , the second equality holds by adding and subtracting $\mathbb{E}[Y_{t^*-1}(0)|D=1]$, and the third equality holds by replacing the potential outcomes in the first term with their observed counterparts. In Eq. (B.13), the first term is directly identified (it is the observed path of outcomes for treated locations) while the second term is not directly identified, and we consider it in detail below. As a preliminary step, notice that

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=0] = \tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=0]$$

Since, $\Delta^{(t^*-1,t)}Y_{it}(0)$ and $\Delta^{(t^*-1,t)}I_{it}(0)$ are both observed outcomes for locations in the untreated group, this implies that $\tilde{\tau}_t$ and α are identified and can be recovered from the regression of $\Delta^{(t^*-1,t)}Y_{it}$ on $\Delta^{(t^*-1,t)}I_{it}$ among the untreated group. Next,

$$\begin{aligned}\mathbb{E}[\Delta^{(t^*-1,t)}Y_{it}(0)|D=1] &= \tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_{it}(0)|D=1] \\ &= \tilde{\tau}_t + \alpha \left(\mathbb{E} \left[\frac{D}{\mathbb{E}[D]} \Delta^{(t^*-1,t)}I_t - \omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^c(\mathcal{F}_{t^*-1})) \right] \right)\end{aligned}$$

and all terms are identified. Plugging this expression into Eq. (B.13) implies the result.

Next, we consider the bias coming from using standard DID. This is given by

$$\begin{aligned}\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=0] - (\mathbb{E}[Y_t(1) - Y_t(0)|D=1]) \\ = \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=0] \\ = \alpha(\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=0])\end{aligned}$$

where the second equality holds by adding and subtracting $\mathbb{E}[Y_{t^*-1}(0)|D=1]$ and canceling terms and the last equality holds by the model in Eq. (11).

Finally, we consider the bias coming from DID including actual cases (rather than cases in the absence of the policy) as a covariate. This bias is given by

$$\begin{aligned}\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=1] - (\tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D=1]) - (\mathbb{E}[Y_t(1) - Y_t(0)|D=1]) \\ = \alpha(\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D=1]) \\ = -\alpha ATT_t^I\end{aligned}$$

where the first equality holds by adding and subtracting $\mathbb{E}[Y_{t^*-1}(0)|D=1]$ and from the model in Eq. (11) (and by canceling terms), and the last equality holds by the definition of ATT_t^I after canceling terms. \square

Appendix C. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jeconom.2023.03.009>.

References

- Abadie, Alberto, 2005. Semiparametric difference-in-differences estimators. *Rev. Econom. Stud.* 72 (1), 1–19.
- Acemoglu, Daron, Chernozhukov, Victor, Werning, Iván, Whinston, Michael D, 2021. Optimal targeted lockdowns in a multigroup SIR model. *Am. Econ. Rev.: Insights* 3 (4), 487–502.
- Alemán, Christian, Busch, Christopher, Ludwig, Alexander, Santaaulàlia-Llopis, Raül, 2020. Evaluating the Effectiveness of Policies Against a Pandemic. Working Paper.
- Allcott, Hunt, Boxell, Levi, Conway, Jacob C, Ferguson, Billy A, Gentzkow, Matthew, Goldman, Benny, 2020. What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic? Working Paper.
- Allen, Linda JS, 2008. An introduction to stochastic epidemic models. In: *Mathematical Epidemiology*. Springer, pp. 81–130.
- Allen, Linda JS, 2017. A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infect. Dis. Model.* 2 (2), 128–142.
- Athey, Susan, Imbens, Guido, 2006. Identification and inference in nonlinear difference-in-differences models. *Econometrica* 74 (2), 431–497.
- Bang, Heejung, Robins, James M, 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics* 61 (4), 962–973.
- Bartik, Alexander W, Bertrand, Marianne, Lin, Feng, Rothstein, Jesse, Unrath, Matthew, 2020. Measuring the labor market at the onset of the COVID-19 crisis. *Brook. Pap. Econ. Act.* 2020 (2), 239–268.
- Bendavid, Eran, Oh, Christopher, Bhattacharya, Jay, Ioannidis, John PA, 2021. Assessing mandatory stay-at-home and business closure effects on the spread of COVID-19. *Eur. J. Clin. Invest.* 51 (4), e13484.
- Berry, Christopher R, Fowler, Anthony, Glazer, Tamara, Handel-Meyer, Samantha, MacMillen, Alec, 2021. Evaluating the effects of shelter-in-place policies during the COVID-19 pandemic. *Proc. Natl. Acad. Sci.* 118 (15), e2019706118.
- Blundell, Richard, Costa Dias, Monica, 2009. Alternative approaches to evaluation in empirical microeconomics. *J. Hum. Resour.* 44 (3), 565–640.
- Bonhomme, Stephane, Sauder, Ulrich, 2011. Recovering distributions in difference-in-differences models: A comparison of selective and comprehensive schooling. *Rev. Econ. Stat.* 93 (2), 479–494.
- Borusyak, Kirill, Jaravel, Xavier, Spiess, Jann, 2022. Revisiting Event Study Designs: Robust and Efficient Estimation. Working Paper.
- Brown, Nicholas, Butts, Kyle, 2022. A Unified Framework for Dynamic Treatment Effect Estimation in Interactive Fixed Effect Models. Working Paper.
- Busso, Matias, DiNardo, John, McCrary, Justin, 2014. New evidence on the finite sample properties of propensity score reweighting and matching estimators. *Rev. Econ. Stat.* 96 (5), 885–897.
- Caetano, Carolina, Callaway, Brantly, Payne, Robert, Rodrigues, Hugo, 2022. Difference in Differences with Time-Varying Covariates. Working Paper.
- Callaway, Brantly, Li, Tong, 2023. Evaluating Policies Early in a Pandemic: Bounding Policy Effects with Nonrandomly Missing Data. Working Paper.
- Callaway, Brantly, Sant'Anna, Pedro HC, 2021. Difference-in-differences with multiple time periods. *J. Econometrics* 225 (2), 200–230.
- Chabé-Ferret, Sylvain, 2015. Analysis of the bias of matching and difference-in-difference under alternative earnings and selection processes. *J. Econometrics* 185 (1), 110–123.
- Chattopadhyay, Ambarish, Zubizarreta, Jose R, 2022. On the Implied Weights of Linear Regression for Causal Inference. Working Paper.
- Chernozhukov, Victor, Kasahara, Hiroyuki, Schrimpf, Paul, 2021. Causal impact of masks, policies, behavior on early Covid-19 pandemic in the US. *J. Econometrics* 220 (1), 23–62.
- Chetty, Raj, Friedman, John, Hendren, Nathaniel, Stepner, Michael, 2023. The Economic Impacts of COVID-19: Evidence from a New Public Database Built from Private Sector Data. Working Paper.

- Cho, Sang-Wook, 2020. Quantifying the impact of nonpharmaceutical interventions during the COVID-19 outbreak: The case of Sweden. *Econom. J.* 23 (3), 323–344.
- Courtemanche, Charles, Garuccio, Joseph, Le, Anh, Pinkston, Joshua, Yelowitz, Aaron, 2020. Strong social distancing measures in the United States reduced the COVID-19 growth rate. *Health Aff.* 39 (7), 1237–1246.
- Crump, Richard K, Hotz, V Joseph, Imbens, Guido W, Mitnik, Oscar A, 2009. Dealing with limited overlap in estimation of average treatment effects. *Biometrika* 96 (1), 187–199.
- Dave, Dhaval M, Friedson, Andrew I, Matsuzawa, Kyutaro, McNichols, Drew, Sabia, Joseph J, 2020a. Did the Wisconsin Supreme Court Restart a COVID-19 Epidemic? Evidence from a Natural Experiment. Working Paper.
- Dave, Dhaval, Friedson, Andrew I, Matsuzawa, Kyutaro, Sabia, Joseph J, 2021. When do shelter-in-place orders fight COVID-19 best? Policy heterogeneity across states and adoption time. *Econ. Inq.* 59 (1), 29–52.
- Dave, Dhaval, Friedson, Andrew, Matsuzawa, Kyutaro, Sabia, Joseph J, Safford, Samuel, 2020b. Were urban cowboys enough to control COVID-19? Local shelter-in-place orders and coronavirus case growth. *J. Urban Econ.* 103294.
- de Chaisemartin, Clement, D'Haultfœuille, Xavier, 2020. Two-way fixed effects estimators with heterogeneous treatment effects. *Amer. Econ. Rev.* 110 (9), 2964–2996.
- de Chaisemartin, Clement, D'Haultfœuille, Xavier, 2023. Two-Way Fixed Effects Regressions with Several Treatments. Working Paper.
- Ellison, Glenn, 2020. Implications of Heterogeneous SIR Models for Analyses of COVID-19. Working Paper.
- Fernández-Villaverde, Jesús, Jones, Charlesa I, 2022. Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. *J. Econom. Dynam. Control* 104318.
- Friedson, Andrew I, McNichols, Drew, Sabia, Joseph J, Dave, Dhaval, 2021. Shelter-in-place orders and public health: Evidence from California during the Covid-19 pandemic. *J. Policy Anal. Manage.* 40 (1), 258–283.
- Fullman, Nancy, Bang-Jensen, Jeremiah, Reinke, Grace, Magistro, Beatrice, Castellano, Rachel, Erickson, Megan, Walcott, Rebecca, Dapper, Carolyn, Amano, Kenya, Wilkerson, John, Adolph, Christopher, 2021. State-level social distancing policies in response to COVID-19 in the US. Version 1.145, <http://www.covid19statepolicy.org/>.
- Gapen, Michael, Millar, Jonathan, Blerina, U, Sriram, Pooja, 2020. Assessing the effectiveness of alternative measures to slow the spread of COVID-19 in the United States. *Covid Econ.* 40, 46–75.
- Gardner, John, 2022. Two-Stage Difference in Differences. Working Paper.
- Gauthier, Germain, 2021. On the Use of Two-Way Fixed Effects Models for Policy Evaluation During Pandemics. Working Paper.
- Ghanem, Dalia, Sant'Anna, Pedro HC, Wüthrich, Kaspar, 2022. Selection and Parallel Trends. Working Paper.
- Glaeser, Edward L, Jin, Ginger Z, Leyden, Benjamin T, Luca, Michael, 2021. Learning from deregulation: The asymmetric impact of lockdown and reopening on risky behavior during COVID-19. *J. Reg. Sci.* 61 (4), 696–709.
- Goldsmith-Pinkham, Paul, Hull, Peter, Kolesár, Michal, 2021. On Estimating Multiple Treatment Effects with Regression. Working Paper.
- Goodman-Bacon, Andrew, 2021. Difference-in-differences with variation in treatment timing. *J. Econometrics* 225 (2), 254–277.
- Goodman-Bacon, Andrew, Marcus, Jan, 2020. Using difference-in-differences to identify causal effects of COVID-19 policies. In: *Survey Research Methods*, Vol. 14.2. pp. 153–158.
- Goolsbee, Austan, Syverson, Chad, 2021. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *J. Public Econ.* 193, 104311.
- Griliches, Zvi, 1977. Estimating the returns to schooling: Some econometric problems. *Econometrica* 1–22.
- Gupta, Sumedha, Montenegro, Laura, Nguyen, Thuy, Lozano-Rojas, Felipe, Schmutte, Ian, Simon, Kosali, Weinberg, Bruce A, Wing, Coady, 2020. Effects of social distancing policy on labor market outcomes. *Contemp. Econ. Policy*.
- Haber, Noah A, Clarke-Deelder, Emma, Feller, Avi, Smith, Emily R, Salomon, Joshua A, MacCormack-Gelles, Benjamin, Stone, Elizabeth M, Bolster-Foucault, Clara, Daw, Jamie R, Hatfield, Laura Anne, Fry, Carrie E, Boyer, Christopher B, Ben-Michael, Eli, Joyce, Caroline M, Linas, Beth S, Schmid, Ian, Au, Eric H, Wieten, Sarah E, Jarrett, Brooke, Axfors, Cathrine, Nguyen, Van Thu, Griffin, Beth Ann, Bilinski, Alyssa, Stuart, Elizabeth A, 2022. Problems with evidence assessment in COVID-19 health policy impact evaluation: A systematic review of study design and evidence strength. *BMJ Open* 12 (1), e053820.
- Hájek, Jaroslav, 1971. Discussion of 'an essay on the logical foundations of survey sampling, Part I', by D. Basu. *Found. Stat. Inference* 326.
- Haynes, Kingsley E, Kulkarni, Rajendra, Li, Meng-Hao, Siddique, Abu Bakkar, 2022. Differences in state level impacts of COVID-19 policies. In: *Theory and History in Regional Perspective*. Springer, pp. 415–431.
- Heckman, James, Ichimura, Hidehiko, Todd, Petra, 1997. Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Rev. Econom. Stud.* 64 (4), 605–654.
- Heckman, James, Robb, Richard, 1985. Alternative methods for evaluating the impact of interventions. In: Heckman, James, Singer, Burton (Eds.), *Longitudinal Analysis of Labor Market Data*. Cambridge University Press, Cambridge, pp. 156–246.
- Hirano, Keisuke, Imbens, Guido, Ridder, Geert, 2003. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 71 (4), 1161–1189.
- Hsiang, Solomon, Allen, Daniel, Annan-Phan, Sébastien, Bell, Kendon, Bolliger, Ian, Chong, Trinetta, Druckenmiller, Hannah, Huang, Luna Yue, Hultgren, Andrew, Krasovich, Emma, 2020. The effect of large-scale anti-contagion policies on the COVID-19 pandemic. *Nature* 584 (7820), 262–267.
- Huber, Martin, 2020. Mediation analysis. In: Zimmermann, Klaus F. (Ed.), *Handbook of Labor, Human Resources and Population Economics*. Springer International Publishing, pp. 1–38.
- Imbens, Guido W, Newey, Whitney K, 2009. Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica* 77 (5), 1481–1512.
- Imbens, Guido, Wooldridge, Jeffrey, 2009. Recent developments in the econometrics of program evaluation. *J. Econ. Lit.* 47 (1), 5–86.
- Juranek, Steffen, Zoutman, Floris T, 2021. The effect of non-pharmaceutical interventions on the demand for health care and on mortality: Evidence from COVID-19 in Scandinavia. *J. Popul. Econ.* 34 (4), 1299–1320.
- Kermack, William Ogilvy, McKendrick, Anderson G, 1927. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A* 115 (772), 700–721.
- Kish, Leslie, 1965. Sampling organizations and groups of unequal sizes. *Am. Sociol. Rev.* 564–572.
- Kong, Edward, Prinz, Daniel, 2020. Disentangling policy effects using proxy data: Which shutdown policies affected unemployment during the COVID-19 pandemic?. *J. Public Econ.* 189, 104257.
- Lechner, Michael, 2011. The estimation of causal effects by difference-in-difference methods. *Found. Trends Econom.* 4 (3), 165–224.
- Manski, Charles F, Molinari, Francesca, 2020. Estimating the COVID-19 infection rate: Anatomy of an inference problem. *J. Econometrics*.
- Marx, Philip, Tamer, Elie, Tang, Xun, 2022. Parallel trends and dynamic choices. Working Paper.
- Mitze, Timo, Kosfeld, Reinhold, Rode, Johannes, Wälde, Klaus, 2020. Face masks considerably reduce COVID-19 cases in Germany. *Proc. Natl. Acad. Sci.* 117 (51), 32293–32301.
- Oka, Tatsushi, Wei, Wei, Zhu, Dan, 2021. The effect of human mobility restrictions on the COVID-19 transmission network in China. *PLoS One* 16 (7), e0254403.

- Roth, Jonathan, 2022. Pretest with caution: Event-study estimates after testing for parallel trends. *Am. Econ. Rev.: Insights* 4 (3), 305–322.
- Sant’Anna, Pedro HC, Zhao, Jun, 2020. Doubly robust difference-in-differences estimators. *J. Econometrics* 219 (1), 101–122.
- Sears, James, Villas-Boas, J Miguel, Villas-Boas, Sofia Berto, Villas-Boas, Vasco, 2023. Are we# stayinghome to flatten the curve?. *Am. J. Health Econ.* 9 (1), 71–95.
- Shook-Sa, Bonnie E, Hudgens, Michael G, 2022. Power and sample size for observational studies of point exposure effects. *Biometrics* 78 (1), 388–398.
- Śłoczyński, Tymon, Wooldridge, Jeffrey M, 2018. A general double robustness result for estimating average treatment effects. *Econom. Theory* 34 (1), 112–133.
- Sun, Liyang, Abraham, Sarah, 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econometrics* 225 (2), 175–199.
- Wright, Austin L, Sonin, Konstantin, Driscoll, Jesse, Wilson, Jarnickae, 2020. Poverty and economic dislocation reduce compliance with Covid-19 shelter-in-place protocols. *J. Econ. Behav. Organ.* 180, 544–554.
- Ziedan, Engy, Simon, Kosali I, Wing, Coady, 2020. Effects of State COVID-19 Closure Policy on Non-COVID-19 Health Care Utilization. Working Paper.